10/26/00

e-Commerce and Operations Research in Airline Planning, Marketing, and Distribution

BARRY C. SMITH — *Sabre Inc.*
1 East Kirkwood Blvd.
MD 7390 TSG
Southlake, TX 76092, USA

DIRK P. GÜNTHER — *Sabre Inc.*

B. VENKATESHWARA RAO — *Sabre Inc.*

RICHARD M. RATLIFF — *Sabre Inc.*

Many e-commerce principles were pioneered in the airline industry. These include the first business-to-business electronic information exchange and industry-wide electronic marketplace. This environment provided unprecedented opportunity for operations research (OR) modeling. By the mid-1980s airlines used customer shopping data to calibrate traveler demand and choice models, analyzed multi-channel product distribution strategies with simulation and practiced dynamic pricing through yield management. Airlines continue to derive billions of dollars annually from these and derivative models. The availability of reliable, low-cost communications via the Internet is not only providing new modeling challenges within the airline industry, but it is also providing similar opportunities in other industries.

The computer reservations systems (CRSs) developed for the travel industry were among the earliest examples of e-commerce. Forty years ago American Airlines developed the SABRE computer reservation system to keep track of seat availability on hundreds of flights for thousands of passengers booked all over the world. Airlines have continuously expanded their use of information technology for reservations systems and created an infrastructure that facilitates the application of OR models.

The SABRE system and similar reservations systems, such as Amadeus, Galileo, Shares, and Worldspan, have transformed the travel industry. They have improved product distribution and customer service, and have revolutionized airline marketing through yield management and frequent flyer programs. These reservation systems have evolved into global distribution systems (GDS). Where CRSs helped individual airlines sell and manage their own seats, the GDS consolidates information from many airlines, allowing travel agents, businesses, and individuals to shop in a single electronic marketplace. This marketplace has expanded to include hotel, rental car, and destination products and services.

The growth of reservation systems and the related use of OR in the airline industry have parallels in the evolving world of Internet-based e-commerce. As noted in *Forbes* online [Malik 1999], “Sabre was an early version of an electronic marketplace, fashionably called an infomediary… The similarities between Sabre and the new infomediaries are quite stark. Infomediaries are web-based market makers that leverage the Internet to bring together buyers and sellers in niche markets, charging commissions for the products they help move… Sabre-type infomediaries took in $290 million in 1998. By 2002, … that number will grow to $20 billion." As stated by Phillips [2000], "Global distribution systems brought the airlines to the electronic marketplace a quarter century before the rest of us."

By consolidating product information from many airline suppliers, global distribution systems created a marketplace with close to perfect information regarding schedules and fares. Using the Internet, other industries are converging towards a similar scenario by adopting such functions as low-price search engines, display-ordering-rules logic and real-time technology that the airlines pioneered. The airlines built up their e-commerce infrastructure over 70 years, and the history of that evolution provides insights into current changes in other industries.

History of Airline e-Commerce

In the prejet commercial era, a few hundred passenger airlines carried just a million passengers per year [Petzinger 1996]. In the 1950s, the introduction of commercial jet engines, cabin pressurization, and improved navigational aids revolutionized air travel. As demand increased, airlines developed internal communications and commercial infrastructure to coordinate the activities of staff, aircraft, and passengers.

SITA (Société Internationale de Télécommunications Aéronautiques) and ARINC (Aeronautical Radio Inc.) were among the world's first business-to-business systems. They used Teletype messaging for communicating electronically among the airlines.

SITA and ARINC are still used today. They employ message pass-through capabilities the airlines developed to transmit information on passengers' travel itineraries and payments. Many businesses have the same communication needs that drove the airlines to develop SITA and ARINC, but costs and technical requirements prevented their developing customized systems. This unmet business need made low-cost communications through e-mail one of the first killer apps of the Internet.

Two of the most difficult information-processing problems the airlines faced were (1) keeping track of the number of seats sold on a growing number of flights and (2) communicating remaining seat availability to geographically dispersed reservation agents. The manual solutions developed in the 1930s were difficult to scale up as the airlines grew. For example, the size of reservation offices was limited by the feasible distance between agents and a chalkboard used to track sold-out flights. At one point, some American Airlines agents used binoculars to read flight availability [Copeland, Mason, and McKenney 1995]. Airlines tried to solve the problem using available technology; the solutions tended to be expensive and barely adequate. In 1953, American Airlines and IBM began working together on the problem. After five years, they formed a partnership to develop the first CRS. In 1962, they implemented the SABRE system for American Airlines' production use. Design, development, and implementation took 400 labor years. It was the first real-time business application of computer technology, an automated system with complete passenger records available electronically to any agent connected to the SABRE system. At that time, American Airlines processed an average of 26,000 passenger reservation transactions per day (in 1999, SABRE processed an average of 1.2 million reservations per day), and the average time to file a reservation decreased from 45 minutes to three seconds. This jump in efficiency allowed airlines to handle the growth in demand for air travel that occurred in the 1960s.

Through the 1970s, the Civil Aeronautics Board (CAB) regulated domestic airfares in the United States. Discount fares were not widely available on scheduled carriers except for narrowly defined demographic groups, such as children and students. Cheaper air travel was available only on unscheduled charter carriers. By the mid-1970s, some chartered carriers started to operate quasi-scheduled flights. Relaxing regulations, the CAB allowed these low-cost carriers to begin competing in the scheduled airline market. Their high costs put the scheduled carriers at a disadvantage. They realized that every empty seat flown was produced at almost zero incremental cost and selling otherwise empty seats at low prices could increase revenue and profits. Initially, they restricted discount fares heavily in an attempt to segment the market and avoid diluting the existing business-fare revenue. Today they still employ some of the initial restrictions, such as advance purchase, round trip, and Saturday night stay.

To fill otherwise empty seats, the airlines increased the number and diversity of discount fares. Airlines developed yield management to fine-tune their increasingly complex pricing structures by tracking reservations and selectively turning on and off the availability of low fares. Yield management applies operations research models to determine the appropriate amount of space to save for late-booking high-revenue customers [Belobaba 1987]. The combination of pricing restrictions and yield-management controls allowed carriers to offer very low fares with little risk of diluting revenues. This was critical for airlines serving both high-revenue business customers and price-sensitive discretionary customers.

Following US airline deregulation in 1978, airlines that controlled their inventory with yield management generated much higher profits than those that did not. A notable illustration is the experience of PeopleExpress. It started service in 1981 and within five years became the fifth largest airline in the United States. At that time, PeopleExpress was the fastest growing company in US history. It could offer low fares because of its no-frills service and nonunionized labor. It had load factors of about 75 percent (75 percent of its seats were occupied) while the rest of the industry had load factors around 55 percent. In January 1985, American Airlines launched a discount pricing campaign, "Ultimate Super Savers," directed at filling the otherwise empty seats. It offered full service at discount fares priced at or below low-cost carriers. It used purchase restrictions and yield-management-based inventory control to prevent the displacement of future full-fare sales by early-arriving discount-fare demand. The overwhelming demand for discount fares and full service shifted traffic to American Airlines and other carriers using similar practices. PeopleExpress load factors dropped to around 25 percent; the airline was sold to Continental in September 1986. According to Donald Burr, former chairman of PeopleExpress, "We were a vibrant, profitable company from 1981 to 1985, and then we tipped right over into losing $50 million a month...What changed was American's ability to do widespread yield management in every one of our markets…" [Cross 1997, p. 124].

As passenger volumes increased in the early 1970s, the paper-intensive processes travel agents used proved too slow. They needed automation to print itineraries, invoices, and tickets, and to handle accounting. In 1974, an effort called the Joint Industry Computer Reservation System (JICRS) was launched to provide automation to travel agents by creating one reservation system for all airlines. Airline participants included American, United, TWA, Eastern, and Western. The plan was to develop a system to display and book reservations on any airline, accommodating the need of travel agents to offer industrywide travel products and services. JICRS was one of the first industry alliances to undertake e-market efforts.

The JICRS effort was disbanded in 1976, when United Airlines announced plans to install its own proprietary system, Apollo, in travel-agent locations. Airlines that owned and operated CRSs rushed to expand their capabilities (to display information for multiple airlines) and to install their systems at travel agencies. The urgency was based on the value associated with display bias. Agents who subscribed to a given carrier's reservation system were provided with flight information biased towards that carrier. For any travel request, many flight alternatives generally exist, even if the traveler has specific preferences (for example, a narrow range of departure or arrival times). The algorithm the GDS used to identify potential itineraries considers schedule convenience (proximity to desired departure time, number of stops, elapsed time) as well as the carrier(s) providing the service. Flights for the carrier operating the GDS were sorted higher in the list of candidate itineraries. The order in which itineraries are presented on the travel-agent screen has a powerful impact on the traveler’s choice. All things considered equal, the first few choices presented have a higher likelihood of being selected. Copeland, Mason, and McKenney [1995] cited findings that an average of 70 percent of all bookings were made from the first screen display (the first six options) presented to a travel agent. By the mid-80s, the incremental revenue associated with display bias for American Airlines was estimated at over $100 million per year [Copeland, Mason, and McKenney 1995].

GDSs (such as SABRE*)* became the gateway for information about all airlines to the travel agent. Control of travel product distribution by a handful of carriers put many airlines at a disadvantage. In November 1984, the Department of Transportation (DOT) regulated several key functions of reservation systems to eliminate bias in their displays. The rulings required *unbiased* displays of flights; public disclosure of display rules by the GDSs; standardization of the timing of daily fare changes (to prevent a carrier from unfairly showing a special fare in its own reservation system before competitors had the opportunity to load their fares); providing reservation system booking data to all airlines (possibly at a charge); and charging equal booking fees. Ironically, the electronic marketplace that had been energized by the deregulation of US airlines was regulated within five years.

The Internet and Travel Distribution

Prior to the Internet, airlines, GDSs, and travel agents communicated across private networks. Public networks supporting the Internet significantly reduced the cost and complexity of communications between the airlines, GDSs, travel agents, and end consumers. Early Internet reservation applications were limited to booking and payment transactions only. Paper tickets were mailed to the travelers, and they had to make purchases seven days in advance if they were to receive their tickets in time to travel. Electronic ticketing facilitated the growth of Internet travel sales by eliminating the need for paper tickets, reducing the lead time and cost associated with online purchases [IBM 2000]. According to Forrester Research, 9 million households booked travel reservations online in 1999. By 2003, this number is expected to surge to 26 million and another 16 million households will research travel online. On the corporate travel front, Forrester projects a revenue volume of $38 billion in 2003 [Walker et al. 1999].

Internet ticket sales have grown rapidly because distribution methods are tailored to customer needs. Three models dominate the Internet distribution channel: (1) Online travel agents (such as Travelocity.com and Expedia) provide direct GDS access to travelers; (2) Airline web sites provide direct links between airlines and customers; (3) Auction and reverse auction outlets bypass the traditional distribution network and act as mediators between suppliers and customers.

Online travel agents, such as Travelocity.com, allow customers to search for travel itineraries that best suit their personal preferences through price-driven and time-driven searches. Customers can also specify additional preference parameters, such as number of connections and airline. Responding to customer's requests requires more resources in the GDS than responding to travel agent requests.

While the Internet has spawned a variety of multisupplier distribution models, it has also revived the airline sales agent in electronic form. Many carriers use the Internet to attract and retain customers by providing additional services such as real-time flight and airport information, online management of frequent flyers, auctions for popular destinations, and online vacation packages. The Internet is also being used to notify known potential travelers about the reduced fares on distressed inventory (Net SAAver Fares , Web Fares). Carriers are offering special promotions, such as exclusive reduced fares and increased mileage, to encourage bookings through their web sites. Some airlines are also restricting promotions to members of their frequent flyer programs. For example, Delta offers "more for your miles" auctions in which members of their frequent flyer program can bid their accumulated miles on vacation packages [Delta 2000]. Airlines are using operations research tools, such as demand forecasting and optimization, to design sales promotions and to identify cross-selling opportunities, such as rental car, hotel, and other destination interests [Lieberman 2000]. They can use customer-profile data captured from the users' navigational paths while browsing the web (their click streams) and other customer-related data as input to identify promotion opportunities, tailored to individual customers. For example, an airline in contact with someone wanting a seat on a completely booked flight might be able to transfer that demand to another flight leaving around the same time or from a nearby airport by e-mailing an offer (flight, price, restrictions) instead of a refusal. Carriers are also using the Internet to cut distribution costs. For example, America West reports average per ticket distribution costs of $23 for traditional channels versus $6 for direct Internet sales [Schwartz and Zea 1999]. While the industry sold two percent of all tickets through airline web sites in 1999, Southwest Airlines sold 15 percent of its tickets through its own web site [*Air Transport World* 2000].

Outlets such as LastMinuteTravel.com are based on concepts of distressed inventory and dynamic pricing. They markdown unsold inventory close to departure and advertise it on the Internet. Customers shopping at this type of outlet must be willing to adapt their travel arrangements to the market; quality (in this case, in the form of travel time and location convenience) is sacrificed for price. The success of these sites (the profit for auction hosts) depends on the choice of sale timing and price. Early discounting (before an airline realizes the regular demand for its flights completely) may result in revenue cannibalization for the airline. Insufficient markdown of fares might not stimulate enough demand. The sites interact with the airlines' revenue management systems seamlessly to determine when and how many seats to auction [LastMinuteTravel.com 2000].

Auction outlets, such as Priceline.com, collect bids from travelers and match them with product offerings from participating airlines. As the customer must accept (within certain parameters) any product that meets the specified price, auction models cater to travelers who have decided to purchase and who are willing to sacrifice quality for price. The auction host's profit is the difference between a customer's bid and the price of the matched product offering. An auction host may opt to accept losing bids to increase revenue. When evaluating a customer bid, the auction host faces the same basic problem as online agents. Hence, the host can use techniques similar to the ones described for the online agent to find cheap fares. If the auction host has access to demand and fare-availability forecasts, it can use yield management to increase its profit margin.

The Internet is creating opportunities for air travel beyond the traditional scheduled airlines. For example, the Internet is enabling charter operators to compete effectively for business travel [Needleman 2000]. Flightserv.com is an online marketplace for charter-business-jet travel. It plans to allow customers to create and book their own routes. Once a customer books a new route, that route will be offered to other potential travelers over the Internet. If successful, this will improve the efficiency of jet charter operations to levels comparable to those of scheduled airlines.

Decision Support in Airline Planning and Marketing

The e-Commerce infrastructure the airlines developed allows collection and central storage of sales and marketing data. Airlines use this data to drive decision support tools for planning and marketing. Figure 1 illustrates the interaction between the different models used for schedule design and resource allocation. The interested reader will find more detailed information in the references provided.

Airline Planning. Decision support models for airline planning are generally directed at estimating the demand, revenue, and cost of proposed schedules. After studies revealed the large impact of GDS screen presence on customer choices, airlines have tried to increase demand by improving their screen presence by adjusting their schedules [Copeland, Mason, and McKenney 1995]. They make small adjustments to improve screen presence, such as (1) determining which origin-destination markets will be served by through flights as opposed to connections in a hub and (2) reducing (or increasing) scheduled flight duration to improve screen presence (and increase demand) at the expense of on-time dependability. Airlines also have modified the fundamental nature of their product offerings to improve screen presence by (1) making agreements with partner airlines to assign multiple carrier codes to the same flight so that interline connections will be displayed as single-airline connections or one-stop through flights and (2) by creating funnel flights where multiple flights are assigned the same numbers so that connections will be displayed as one-stop through flights. While many considered these practices deceptive, they became industry standards as airlines pursued screen presence.

The tasks of estimating demand and determining appropriate schedule adjustments are daunting. Airlines developed CRS simulation models to estimate screen presence and the resulting market share [Nason 1981]. Airlines that owned GDSs used CRS simulations to determine the display parameters (within regulatory limits) that most favored their schedules. Given these published GDS parameters, other airlines used CRS simulations to improve their screen presence. Market-share models are calibrated using market information data or customer information data that contain booking information. Customer information data consist of protocols (such as keystroke sequences) of randomly selected sales sessions that occur during the booking process. These data are analogous to click-stream data and are used to calibrate discrete customer choice models typically based on multinomial logit or nested logit approaches [Ben-Akiva and Lerman

1985]. Customer information data are, in general, regarded as the best data source for building passenger preference models based on screen presence.

The market-share estimates produced by CRS simulation models are essential parts of airline profitability analyses. Airlines use the market-share estimates in several stages of the schedule-design process to determine the basic schedule structure, service frequency, and timing [Smith, Barlow, and Vinod 1998]. Once an airline establishes the base schedule, it uses a fleet assignment model to assign aircraft types to flights. The objective of this model is to maximize captured revenue minus operating cost while meeting various operational constraints [Abara 1989; Jacobs, Smith, and Johnson 1999]. The primary input to the fleet assignment model for estimating demand and revenue comes from the airline profitability model. Currently the airlines manually develop the base schedules and input them to the fleet assignment model. Several efforts are underway to develop decision support for this initial step of schedule design [Barnhart et al. 2000; Lettovský, Johnson, and Smith 1999].

Display regulations governing the traditional GDSs do not apply to the Internet-based distribution channels. Online travel agents and auction outlets can establish their own display rules. To date, airlines have not taken advantage of the Internet to develop new processes for scheduling flights. In the future, we expect that airlines will collect data corresponding to the Internet-based demand and then design a schedule that best accommodates this demand, or auction space on provisional flights and operate the schedules that will collectively produce the greatest profit. So far, no major carriers have announced such plans.

Other industries that distribute products through Internet channels have product-planning problems similar to those the airlines faced. They need models to track product visibility, to adjust products to the channels, and to estimate the impact on demand and revenue. If privacy issues are resolved [*Business Week* 2000], detailed customer-shopping information will become widely available. Firms will likely use discrete-choice modeling to estimate market share and product demand [Montgomery 2001]. Firms, such as Amazon.com, are using collaborative filtering and similar techniques to personalize web-site content by predicting individual preferences based on the navigational paths of customers considered to have similar tastes [Pescovitz 2000].

Airline Marketing. The most important modeling application in airline marketing is yield management. American Airlines benefited by $1.2 billion over a three-year period from its yield-management practices. Yield management systems use historical flight-demand and passenger-booking data to set reservation availability controls [Smith, Leimkuhler, and Darrow 1992]. Airlines typically use several types of models in these systems. They use forecasting models to predict future demand and cancellations and base three types of decisions on these forecasts: flight overbooking, discount-fare management, and itinerary control. They use overbooking models to determine the maximum number of reservations (generally more than physical capacity) to accept. They use fare-management models to determine when to stop selling discount fares to avoid losing later-booking, higher-valued sales. They use traffic models to determine the best mix of long-haul (usually connecting) and short-haul (local) passengers. The contention for seats caused by multiple origin-destinations flowing over the same flights, combined with multiple fares for each origin-destination pair, and the uncertainty of underlying demand requires a stochastic optimization approach to yield management.

Airlines generate additional revenue by carrying freight on passenger flights and by operating separate freighters; in 1998 Lufthansa made $1.331 billion (17 percent of total revenue) by shipping air cargo [*Airline Business* 1999]. Following the success of yield management practices in passenger business, airlines devised similar techniques to improve the profitability in air cargo business; Cathay Pacific Cargo estimated that its revenues increased 3.8 percent in the fourth quarter of 1996 because of yield management [*Payload Asia* 1997]. Sabre developed a cargo-routing guide that finds alternative routes that meet the requested service and shipment characteristics. Airlines use the cargo-routing algorithm when they compute the yield-management controls and when they respond to online customer requests. When airlines set yield management controls, they use the routing model to identify alternative routings for each market and service level. The yield management controls are computed by considering the combined capacities of all routing alternatives [Rao 2000]. When responding to customer requests, airlines use the routing algorithm to find the cheapest feasible route. The routing problem is solved by using a shortest path algorithm with side constraints [Günther et al. 1999]. The algorithm generates a hierarchical decomposition of the flight network. A modified bidirectional Dijkstra-type algorithm uses the decomposed network to respond to online routing requests.

Internet distribution has changed the yield management problem drastically. As airlines sell their excess inventory through Internet auctions and other outlets, they are extending yield-management practices to include dynamic pricing and focused targeting of discounted fares and promotions to individuals. The new channels enable airlines to capture previously unobserved portions of the price-demand curve: the demand associated with low fares (auction and unpublished fares). This is changing the airline pricing structure from oligopoly (whereby airlines publish their fares via a single forum, the Airline Tariff Publishing Company) to retail (where each airline sets up its own fare products). This focused marketing is affecting the forecasting and optimization algorithms airlines use to compute yield-management controls, because these fares are lower in value and in cancellation risk.

Airline decision support systems were extended to other travel-related industries. Sabre and SNCF (the French national railroad, Societé Nationale des Chemins de Fer Française) jointly developed a schedule-planning system to determine fleet mix and configuration, service frequency, timing, and railcar routing for high-speed rail service [Ben-Khedher et al. 1998]. Derivatives of the airline profitability model, the fleet assignment model, and the fleet routing model have all been applied to the passenger rail business [Ciancimino et al. 1999]. Unlike airlines, SNCF publishes schedules for only 80 percent of its high-speed capacity. It schedules the remaining 20 percent during the final

14 days prior to departure based on the strength of advanced booking. This flexibility allows SNCF to closely match supply to demand. Sabre and SNCF jointly developed a model to assign and route this incremental capacity in the network to maximize profit. Researchers have advocated this procedure of adjusting capacity close to departure for many years [Berge and Hopperstad 1993], and airlines such as Continental are now attempting to make it a practice [Pastor 1999].

Sabre has also helped Amtrak, SNCF high-speed trains, and Eurostar (the channel tunnel high-speed service) develop yield-management processes. The hotel and cruise industries have also adopted yield management systems, using forecasting, overbooking, and price management. Hotels can benefit by determining the best mix of one-night versus multiple-night reservations [Bitran and Mondschein 1995]. The problem is similar to that of determining the best mix of short-haul versus long-haul reservations for airlines, so hotels can rely on corresponding models. Since hotels and cruise lines can have many different room (or cabin) types, customer upgrades add another dimension of optimization complexity.

While these applications of scheduling and yield management are closely related to the airline experience, the Internet has made yield management modeling practical for a much wider range of industries, such as broadcasting, retail, manufacturing, and power generation [Cross 1997; Geraghty and Johnson 1997; PROS Revenue Management; TALUS Solutions]. Karaesmen and van Ryzin [1998] describe the analogies between yield management in travel and transportation and inventory and production theory. The random yield in inventory and production is similar to the number of reservations remaining after cancellations. The allocation of manufacturing capacity and inventory to various product demands is identical to determining availability controls for various travel itineraries. Despite these similarities, there is a major difference. The capacity in travel and transportation is fixed, while the capacity in manufacturing tends to be flexible. Goodwin et al. [1999] extended the concepts of yield management to business-to-business exchanges in the manufacturing environment. They developed a distributed decision support framework to match the capacity and demand across multiple suppliers and consumers using the Internet search engines and yield-management models. Priceline.com has extended the potential application of yield management even further by creating a business model that can support yield management of home mortgages, new car sales, long-distance-calling minutes, and groceries. While it is not clear to what extent Priceline.com practices yield management, the potential clearly exists.

Without the Internet, none of these industries could have developed the infrastructure or critical mass of consumers to support this type of flexible marketplace. The wide applicability of yield management and its value to both consumers and suppliers may make it one of the next killer apps of the Internet.

Modeling in Travel Distribution

A large amount of processing resources within the GDS is required to respond to shopping requests, that is, finding the best flight itineraries and fares based on travel agent or consumer requests. For example, about 50 percent of SABRE's CPU resources is dedicated to this function.

SABRE, like the other major GDSs, stores data in a schedule-oriented way. While this provides for very efficient schedule and itinerary searches, it does not promote efficient price-driven searches. The volume of price-driven requests, which initially increased following airline deregulation, has skyrocketed following the introduction of the Internet channels directed at bargain-seeking consumers. The number of price-driven requests the SABRE system handles has approximately doubled each year for the past three years. The search strategies used in SABRE that evolved fairly slowly between 1980 and 1997 required major retooling to keep up with the requirements of Travelocity.com and other growing channels.

In 1984, Sabre developed the Bargain Finder℠, which takes a customer's itinerary as input, performs an extensive fare search, checks fare rules, and returns the cheapest fare class combination for that itinerary. Despite its limited functionality, Bargain Finder proved very valuable to travel agents because fare rules are complex and hard to check. In 1993, Sabre developed an enhanced version called Bargain Finder Plus℠, which considers alternative fare classes on the given itinerary and alternative flights. It uses a two-step heuristic. In the first step, it uses local fares to select a set of alternative flights for each portion of the itinerary. In the second step, it generates and ranks a subset of all combinations of the previously selected flights. It displays the best itineraries.

Bargain Finder Plus was very well received by its users and quickly became a widely used tool in SABRE. However, because the two-phase heuristic performed extensive, complex fare-rule checking to establish valid fares for single flights and itineraries, running Bargain Finder Plus took more of SABRE's CPU resources than anticipated. In addition many of the itineraries Bargain Finder Plus produced had very similar characteristics. Often, alternatives differed by only a few dollars or by minor differences in departure time, leaving the user with few real alternatives.

In 1997, Walker [1997] proposed a dynamic programming algorithm using marginal utilities to improve Bargain Finder Plus performance. In the revised model, flight utilities are based on such characteristics as number of stops, airline, total travel time, and timeliness. The algorithm builds itineraries with maximum marginal utility and produces itineraries with the desired diversity. Moreover, it greatly reduced CPU usage, saving Sabre about $6 million per year.

With the advent of Travelocity.com and related sites, the need for tools like Bargain Finder Plus increased dramatically. Prior to Travelocity.com, travel agents were the only users of Bargain Finder Plus. They typically used Bargain Finder Plus to perform fine-tuning at the end of the booking process. Customers booking over the Internet more frequently rely on low-fare search tools as their main information sources, so efficient algorithms that support low-fare searches are essential for Travelocity.com and other online agents. An ideal search algorithm uses a customer-specific preference profile to construct a utility function addressing service and fare, and efficiently finds itineraries that maximize this utility function. Since the customers' utility functions are unknown (often even to themselves), the goal of the search algorithm is to produce a diverse set of options that has a high likelihood of containing an acceptable (if not the optimal) itinerary.

From a modeling perspective, Travelocity.com and related sites face a k-shortest (least-cost) path problem with side constraints. They need to find a feasible set of routes that best fits the selected utility function. Problem size, complex fare rules, and dynamic data make the instances that arise in the context of low-fare search particularly hard to solve. For example, a customer wants to book a round trip from New York City (NYC) to Los Angeles (LAX). For any given departure date, there are over 2,000 fare-flight combinations for each way of the round trip on 12 airlines providing service. That number increases if we include interline connections, where a passenger changes airlines along the way. Moreover, the fares filed for an origin-destination pair can be undercut by using other stations as connection points. With six additional connection points, the network will contain over 5 million paths from NYC to LAX. Each path represents a combination of fare and itinerary. This example shows the importance of eliminating dominated paths early in the search process, especially since the problem has to be solved multiple times per second. Discount fares are restricted by complex fare rules that make it hard to identify feasible paths in the network. Some rules specify valid travel dates, advance purchase requirements, or minimum stay restrictions. Others, so called *combinability* rules, restrict fare combinations. The first type of rule can be checked when the fare is selected, while the latter must be checked when each path is constructed.

Sabre developed a shortest-path algorithm that uses data aggregation and bounding techniques to limit the search space. The algorithm builds a search tree using the departure station as root node. Each node of the tree is associated with a partial itinerary. In each iteration, the algorithm uses data aggregation to establish a lower bound on completing a partial itinerary. At this stage, it evaluates fare rules that restrict fare combinations to make bounds as tight as possible [Walker and Zhang 1998]. It considers the actual node costs plus the lower bound on completion cost when expanding the search tree. A shortest path approach to the low-fare search problem gives users such information as the cheapest possible fare to complete the requested route along with conditions that need to be satisfied to get this fare. Other modes of transportation, lodging, and rental car options can be included in the search as well. Hence, the shortest path algorithm can be used to build customized vacation packages. Moreover, it can be used to choose from a pool of prepackaged vacations.

Search engines similar to Sabre's Bargain Finder Plus operate using the data available in other GDSs; Best Buy Quote, Power Pricing, and Seven Day Search operate on Apollo/Galileo, Worldspan, and Amadeus, respectively. Unlike these search engines, ITA Software's pricing tool bypasses the GDSs and works directly with data on the airlines' published and private fares and rules. The ITA search engine sifts efficiently through over a billion fare and flight combinations for each travel request [Bly 2000]. The specific techniques employed in these search engines are not published in the literature due to the confidential nature of the subject.

Travel Distribution Neutrality and Regulation

Many early travel web sites (including Sabre's) were at least partially owned by airlines, and new ones are planned. For example, a consortium of major carriers is collaborating on a new travel web site, Orbitz [*CNNfn* 1999]. Unless the carriers clearly state that the site is airline-owned, Orbitz customers could be misled into thinking that they are getting a complete and unbiased range of alternatives. The US Department of Justice is investigating whether Orbitz is anti-competitive [Dobbyn 2000]. Regardless of the outcome of the investigation, the self-regulation driven by market forces will probably have a large impact on the neutrality of airline-owned reservation sites. These market forces will ultimately promote the neutrality of sites operated by multiple suppliers.

Even in genuinely neutral infomediaries, bias in some form is unavoidable for technical and commercial reasons. First, any system will require rules that determine what to display. The factors affecting the display ordering of competing alternatives can be weighted. For example, connecting flights can be given less weight than direct services, interline connections can be penalized compared to single-carrier connections, and multistop flights can be disadvantaged compared to nonstop flights. By analyzing the displays that result from different weighting combinations, one can develop selection rules that favor one carrier over another. Second, common business practices, such as suppliers (airlines) offering volume-based commissions, create powerful incentives for web sites to promote certain carriers over others.

The airlines' rush in the 1970's to penetrate travel agencies with new automation tools has a clear parallel to the Internet economy. While airlines developed automated distribution tools for their own benefit, Internet distribution is an industry in its own right. Portals and search engines derive their value by providing access to many products and services. It is not uncommon for portals to contain exclusive links to service providers and e-retailers. For example, MSN.com links to a single travel-service provider, Expedia; and America Online has agreed that Travelocity.com will be its exclusive travel-service link. Priceline.com gives Delta a priority treatment [Sileo 2000]. Unlike other suppliers, Priceline.com uses only highest qualifying fares to fill requests on Delta. Also, Delta is entitled to revenue sharing, has the right to approve new airline participants, and may prevent Priceline.com from serving some markets.

The technical and commercial issues airlines and GDSs face are those any infomediary displaying goods and services of multiple vendors (such as mortgages, insurance, and personal computers) must face. Generally infomediaries must make money. Suppliers are paying for two things: (1) inclusion on the web site and (2) prominent display. Display bias invites government regulation. In assessing bias, one must decide whether users view the web site as an extension of the supplier. Customers expect a company's own web site to be biased. For example, Dell or Micron wants to sell its own personal computers. However, if customers view the web site as a "public utility" (that is, as a neutral site), bias can be insidious. The airlines have learned to tailor their products (flight schedules and airfares) to different distribution channels (originally the major GDS vendors and now the Internet) to maximize their screen presence and sales effectiveness. Other companies doing business on the Internet must consider the screen presence of their products (relative to those of their competitors) on the leading web sites. Companies can use decision support tools to assess product visibility and channel presence, so as with the airlines, suppliers that fail to do so will be at a competitive disadvantage.

Operations Research in the e-Commerce Environment

The success of airline applications and the communications, data, and control afforded by the Internet are encouraging other industries to apply decision support models. *The Wall Street Journal* summarized the situation: “Eventually, many suppliers are likely to use the Web’s fine-tuned interactivity to perfect yield management strategies similar to the way airline tickets are priced today, slashing prices to avoid surplus inventory or to quickly respond to changes in customer preferences.” [Bank 1999].

The evolution of OR models for the airlines provides valuable guidance to the application of modeling to new industries facilitated by the Internet. Both depend on (1) the availability of data to support planning and marketing decisions; (2) suppliers’ control of price and product availability; and (3) Dynamic “storefronts” provided by infomediaries who control product visibility. Since the airlines started developing their models, however, the pace of change has quickened. Potential applications for forecasting, pricing, yield management, scheduling, and resource allocation arise almost daily. Development time has been reduced. This reduction in the development time will put OR modeling on the critical path with greater frequency. As a result we must develop models faster. We must use high-level modeling languages and tools for this environment. The pressure to develop applications rapidly and the high expectations associated with them are raising the risks and the rewards for OR modeling. Even the most successful models don't get it right every time, and many don't get it right the first time. An optimization-related glitch was estimated to have cost American Airlines about $50 million in 1988 [Cross 1997]. We are seeing similar high-profile problems in some other applications. Because it underestimated demand, the Encyclopedia Britannica web site failed to serve most of the customer requests during its launch. Because it greatly overestimated demand, Williams-Sonoma, an upscale kitchenware and home-furnishing retailer, was left with 40 percent higher post-holiday inventory levels than in the previous year [Lee 2000]. In this evolving environment, OR practitioners need to respond to extreme pressures to cut the time to market but ensure that they take enough time to maintain reliability and credibility. This is particularly true when decisions involve multiple components, such as forecasting, optimization, and delivery.

Some modeling issues have not changed. Data availability and quality is often the most constraining aspect of an OR project. The data associated with new e-commerce applications may be cleaner than those from legacy systems; however, these data are often used along with data from the existing systems. Combining and cross-referencing data from multiple applications often highlights errors in each data source and inconsistencies between sources. It seems unlikely that data from the Internet will take less time and effort to verify. Second, successful modeling approaches should include provisions for measuring their benefits. Unfortunately, the modeling task to measure benefits is often as challenging as the task of developing the model to be measured. Performance measurement can have major impact. For example, Sabre has worked on major reservation system developments justified solely by the incremental revenues obtained from yield management. Performance measurement will continue to be essential.

Finally, as business processes change rapidly, models will likely have shorter life cycles. This may appear problematic, because businesses have less time to benefit from successful models. But they may also benefit. It has always been difficult to maintain the core mathematical models in decision support systems over the long term. Ellis Johnson coined the term *model petrification* to refer to the organization's loss of understanding of models after the original development team departs. As responsibility for the models passes from team to team, the organization gradually loses insight into the models. Often it can no longer enhance or adapt the models to changing business requirements. Models sometimes stop working or worse, begin to give inappropriate results that may go unnoticed. Working at Internet-speed, modelers may be in place longer than their models. Practitioners may have more career opportunities with jobs that emphasize development over maintenance, and model performance and impact should improve.

Acknowledgements

We are grateful to Arthur Geoffrion, Ramayya Krishnan, and four anonymous referees for their thoughtful and extensive comments that led to major improvements of the manuscript. We also thank Michael Clarke, Vinay Dube, Bob Offutt, Abdoul Karim Sylla, Pat Trapp, and Alan Walker for their valuable comments and discussions.

List of URLs

ARINC, www.arinc.com

Expedia, www.expedia.com

Flightserv.com, www.flightserv.com

ITA, www.itasoftware.com/ourtechnology.html

LastMinuteTravel.com, www.lastminutetravel.com

Net SAAver Fares,
www.im.aa.com/American?BV_Operation=Dyn_AAPage&referer=index.html&form%25referrer_site=None

Priceline.com, www.priceline.com

PROS Revenue Management, www.prosrm.com

SITA, www.sita.com

TALUS Solutions, www.talussolutions.com

Travelocity.com, www.travelocity.com

Web Fares, www.delta-air.com/travel/sp_offers/web_fares/index.jsp

References

Abara, Jeph 1989, “Applying integer linear programming to the fleet assignment problem,” *Interfaces*, Vol. 19, No. 4 (July-August), pp. 20-28.

Air Transport World 2000, "News briefs: Big will get bigger in online travel," Vol. 37, No. 3 (March), pp. 11-14.

Airline Business 1999, “Special report: Cargo,” Vol.15, No. 11 (November), pp. 53-62.

Bank, David 1999, “E-commerce (a special report) --- a new model – a site-eat-site world: Disappearing profit margins have retailers fretting – and consumers rejoicing,” *The Wall Street Journal*, July 12, 2000.

Barnhart, Cynthia; Armacost, Andrew P.; Farahat, Amr; and Lohatepanont, Monoj 2000, “Airline schedule planning,” Paper presented at the INFORMS Spring 2000 Meeting, Salt Lake City, Utah, USA.

Belobaba, Peter P. 1987, “Airline yield management: An overview of seat inventory control,” *Transportation Science*, Vol. 21, No. 2 (May), pp. 63-73.

Ben-Akiva, Moshe and Lerman, Steven R. 1985, *Discrete Choice Analysis: Theory and Application to Travel Demand*, The MIT Press, Cambridge, Massachusetts, USA.

Ben-Khedher, Nejib; Kintanar, Josephine; Queille, Cecile; and Stripling, William 1998, “Schedule optimization at SNCF: From conception to day of departure,” *Interfaces*, Vol. 28, No. 1 (January-February), pp. 6-23.

Berge, Matthew E. and Hopperstad, Craig A. 1993, “Demand driven dispatch: A method for dynamic aircraft capacity assignment, models and algorithms,” *Operations Research*, Vol. 41, No. 1 (January-February), pp. 153-168.

Bitran, Gabriel R. and Mondschein, SusanaV. 1995, "An application of yield management to the hotel industry considering multiple day stays," *Operations Research*, Vol. 43, No. 3 (May-June), pp. 427-443.

Bly, Laura 2000, “Upstart airfare site beats the big boys,” *USA Today.com*, April 21. Retrieved on 23 October 2000 from www.usatoday.com/life/travel/leisure/2000/ltl227.htm

Business Week 2000, “It’s time for rules in wonderland,” March 20, pp. 83-96.

Ciancimino, A.; Inzerillo, G.; Lucidi, L.; and Palagi, L. 1999, "A mathematical programming approach for the solution of the railway yield management problem," *Transportation Science*, Vol. 33, No. 2 (May), pp. 168-181.

CNNfn 1999, “Airlines join in web portal: United, Delta, Northwest, Continental to offer joint Internet travel site in 2000,” November 09. Retrieved on 23 October 2000 from cnnfn.cnn.com/1999/11/09/travelcenter/airweb/index.htm

Copeland, Duncan G.; Mason, Richard O.; and McKenney, James L. 1995, “Sabre: The development of information-based competence and execution of information-based competition,” *IEEE Annals of the History of Computing*, Vol. 17, No. 3 (Fall), pp. 30-55.

Cross, Robert G. 1997, *Revenue Management: Hard-Core Tactics for Market Domination*, Broadway Books, New York, USA.

Delta 2000, “More for your miles – Auction showcases Delta destinations and amazing adventures,” Press release, May 22. Retrieved on 23 October 2000 from www.delta-air.com/home/press/archiveindex.jsp?tactype=Archived

Dobbyn, Tim 2000, "U.S. Justice Department probes airline ticket web site," *Reuters Business News*, May 18. Retrieved on 23 October 2000 from abcnews.go.com/news/sections/business/dailynews/reuters20000518_3647.html

Geraghty, M. K. and Johnson, Ernest 1997, "Revenue management saves National Car Rental," *Interfaces*, Vol. 27, No. 1 (January-February), pp. 107-127.

Goodwin, Richard; Keskinocak, Pinar; Murthy, Sesh; Wu, Fredrick; and Akkiraju, Rama 1999, "Intelligent decision support for the e-supply chain," White paper, IBM T. J. Watson Research Center, Yorktown Heights, New York, USA. Retrieved on 23 October 2000 from www.research.ibm.com/CoopDS/aiec99.PDF

Günther, Dirk; Johnson, Ellis L.; Saad, Nader; and Vande Vate, John H. 1999, "An online routing algorithm," Paper presented at the AGIFORS Cargo Study Group, New Orleans, LA, USA.

IBM 2000, "Airline electronic ticketing solutions". Retrieved on 23 October 2000 from houns54.clearlake.ibm.com/solutions/travel/trapub.nsf/SearchResults/8625680F007008ED862568F1007B5797

Jacobs, Tim L.; Smith, Barry C.; and Johnson, Ellis L. 1999, "O&D FAM: Incorporating passenger flows into the fleeting process," *AGIFORS Symposium Proceedings*, Vol. 39, New Orleans, Louisiana, USA.

Karaesmen, Itir and van Ryzin, Garrett 1998, "Overbooking with substitutable inventory classes," Working paper, Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA.

LastMinuteTravel.com 2000, "LastMinuteTravel.com and PROS Revenue Management build the runway to the information super highway," Press release, February 03. Retrieved on 23 October 2000 from
www.lastminutetravel.com/media/pressreleases/feb-3-2000.html

Lee, Louise 2000, "An e-commerce cautionary tale: As Williams-Sonoma's site faltered, inventory piled up," *Business Week*, March 20, 2000, pp.46-48.

Lettovský, Ladislav; Johnson, Ellis L.; and Smith, Barry C. 1999, "Schedule generation model," *AGIFORS Symposium Proceedings*, Vol. 39, New Orleans, Louisiana, USA.

Lieberman, Warren 2000, "The integration of pricing and yield management," Paper presented at the INFORMS Spring 2000 Meeting, Salt Lake City, Utah, USA.

Malik, Om 1999, “Startup muse,” *Forbes.com*, August 18. Retrieved on 23 October 2000 from www.forbes.com/1999/08/18/feat.htm

Montgomery, Alan L. 2001, “Applying quantitative marketing techniques to the Internet,” *Interfaces*, Vol. 31, No. 2 (March-April).

Nason, Scott D. 1981, “The airline preference problem: An application of disaggregate logic,” *AGIFORS Symposium Proceedings*, Vol. 21, Santa Barbara, California, USA.

Needleman, Rafe 2000, “Gotta fly,” *Redherring.com*, March 30. Retrieved on 23 October 2000 from www.redherring.com/cod/2000/0330.html

Pastor, Judy 1999, “Demand driven dispatch: A stochastic approach,” *AGIFORS Symposium Proceedings*, Vol. 39, New Orleans, Louisiana, USA.

Payload Asia 1997, “Focus - Cathay Pacific adopts yield management to improve cargo revenues,” Vol. 13, No. 4 (April), pp. 32-33.

Pescovitz, David 2000, “Accounting for taste,” *Scientific American,* Vol. 282, No. 6 (June), page 44.

Petzinger, Thomas, Jr. 1996, *Hard Landing: The Epic Contest for the Power and Profits that Plunged the Airlines into Chaos*, Random House, New York, USA.

Phillips, Robert L. 2000, "E-commerce: Revenue management at the speed of light," White paper, TALUS Solutions, Waterstone, Suite 300, 4751 Best Road, Atlanta, Georgia, USA

Rao, B. Venkateshwara 2000, "An origin-destination model for cargo revenue-mix optimization," Paper presented at the AGIFORS Cargo Study Group, Louisville, Kentucky, USA

Schwartz, Nancy and Zea, Michael 1999, "Surfing for value," *Airline Business*, Vol. 15, No. 7 (July), pp. 68-70.

Sileo, Lorraine 2000, "Priceline.com: Unproven business model, but Delta is on board," *The Insighter*, Vol. 1, No. 14 (January 22). Retrieved on 23 October 2000 from http://ecommerce.internet.com/opinions/article/0,,3551_111561_2,00.html.

Smith, Barry C.; Leimkuhler, John F.; and Darrow, Ross M. 1992, "Yield management at American Airlines," *Interfaces*, Vol. 22, No. 1 (January-February), pp. 8-31.

Smith, Barry C.; Barlow, Jim; and Vinod, Ben 1998, "Airline planning and marketing decision support: A review of current practices and future trends," *Handbook of Airline Marketing*, eds. G.F. Butler and M.R. Keller, McGraw Hill, New York, USA, pp. 117-130.

Walker, Alan 1997, "TripSearch Phase 1 using marginal utility," Internal report, Sabre, Inc.

Walker, Alan and Zhang, Yanjun 1998, "A new low-fare search engine", Internal report, Sabre, Inc.

Walker, Joshua; Schadler, Ted; Worthen, Ben; Ciardelli, J. Amanda; and Smith, Stephanie 1999, " Travel's Internet backbone," Report, Forrester Research, Inc., December.

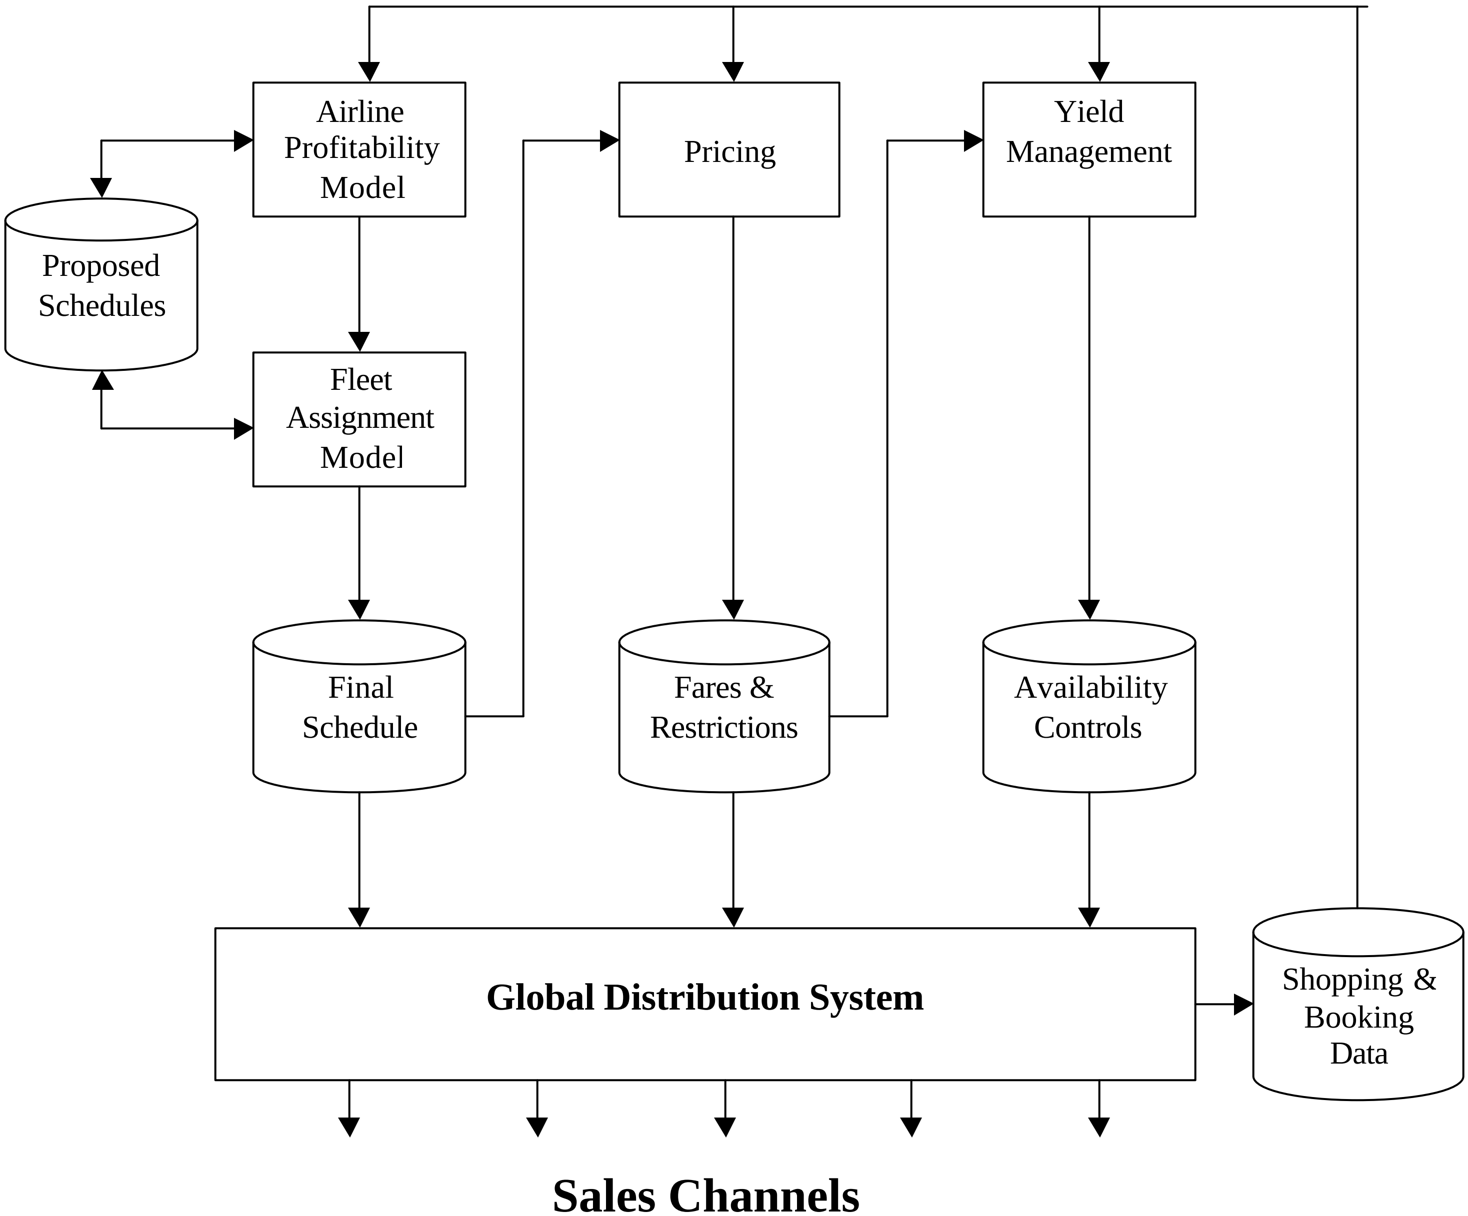

Figure 1: Data collected and stored using the e-commerce infrastructure drive the airline scheduling and yield management. Airlines use the sales and marketing data to build "optimal" schedules. They also use the data to set "optimal" prices and yield management controls.