

Demand Driven Dispatch



Contents

Introduction	1
The Problem	2
The Solution - Demand Scheduling	6
Boeing's D ³ Simulator	9
The Payoff	13
D ³ in Operation	14
Case Study	15

Introduction

One of the major challenges facing the airline industry today is the assignment of airplane capacity to flight schedules to meet fluctuating market needs.

The idea of demand-based scheduling – assigning the right size airplane to each departure – has been a tantalizing but elusive concept for years.

With the advent of common-crew-rated airplane families, notably the Boeing 737 family, the concept can more realistically become a reality.

Boeing calls this concept Demand Driven Dispatch or D³. By coordinating capacity assignment with yield management, the airline can achieve a better match between the supply of seats and passenger demand.

The result is improved contribution to operating profit.

Boeing has developed the D³ Simulator to examine the concept and to explore the benefits of Demand Driven Dispatch.

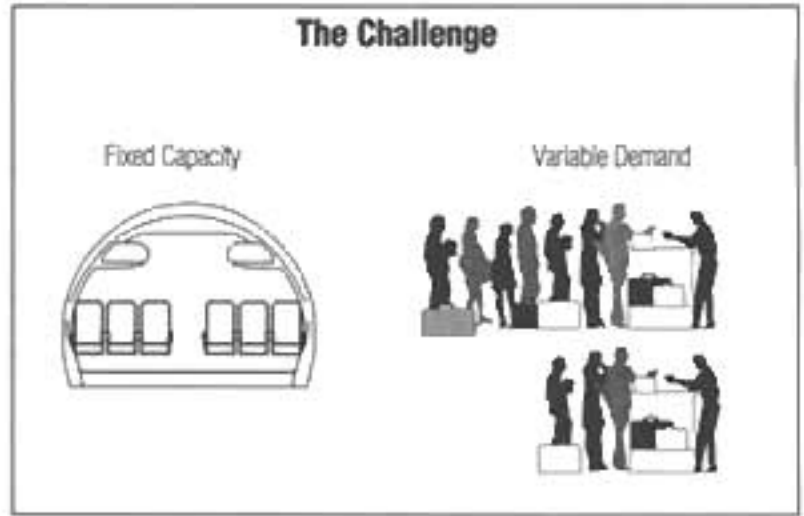
This brochure discusses the theory and benefits of D³ and summarizes how the Boeing model works.



The Problem

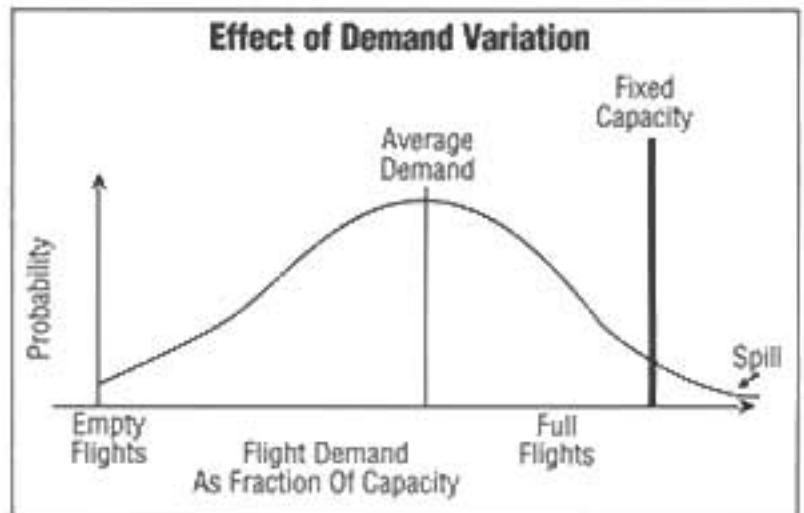
Assigning airplane capacity to flight schedules is a challenge because demand varies widely within each market. While the majority of flights will depart with average loads, some could be completely full and others disturbingly empty.

The variation reflects seasonality, passenger preference for day of week and time of day, passenger response to competitor actions, plus the inherent random fluctuation of the passenger population, day by day, route by route, departure by departure.




When airplane capacity is fixed and demand is lower than expected, the result is low load factor, wasted resources and unrecoverable operating costs.

Greater than expected demand for a fixed-capacity flight results in lost revenue from passenger turnaway (spill), and the associated loss of passenger goodwill.



Today the airline industry operates at about 65 percent passenger load factor. Obviously, a 65 percent load translates to 35 percent of the seats going out empty, with significant potential revenue lost forever.

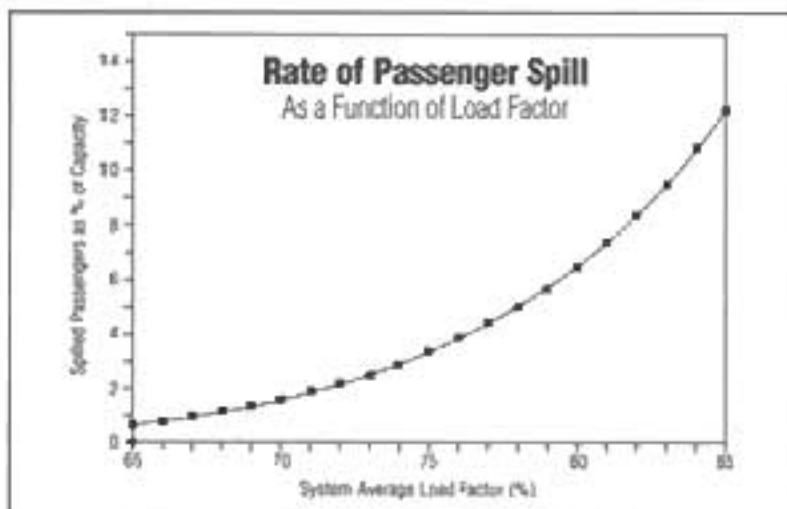


The Problem (cont.)

Reducing capacity is not the solution.

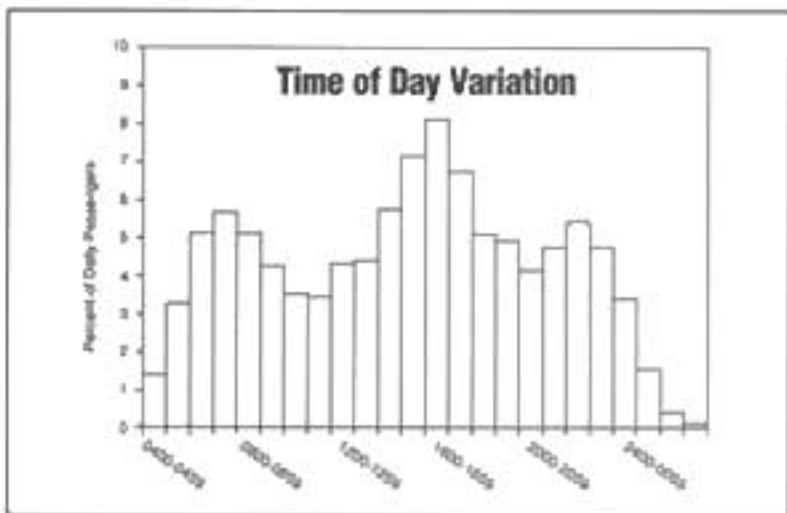

From the traveler's perspective, a substantial improvement in system load factor could translate to a sharp increase in the probability of being turned away when calling for a reservation... leading to loss of the passenger's future business... leading to erosion of the airline's market share.

The problem is not a simple matter of excess capacity but rather it is a problem of not having the **right capacity at the right place at the right time**.



Some of the factors which cause demand to vary are explainable by traveler preferences for certain seasons, days of the week, and times of the day.

Airline pricing strategies and yield management systems have helped to smooth this type of variation. The earliest task of yield managers was to partition the capacity of each flight appropriately by fare class, to accommodate the full fare traveler at the demand peaks and generate discretionary traffic to fill the valleys. In this respect, yield management systems have improved airplane revenue significantly.

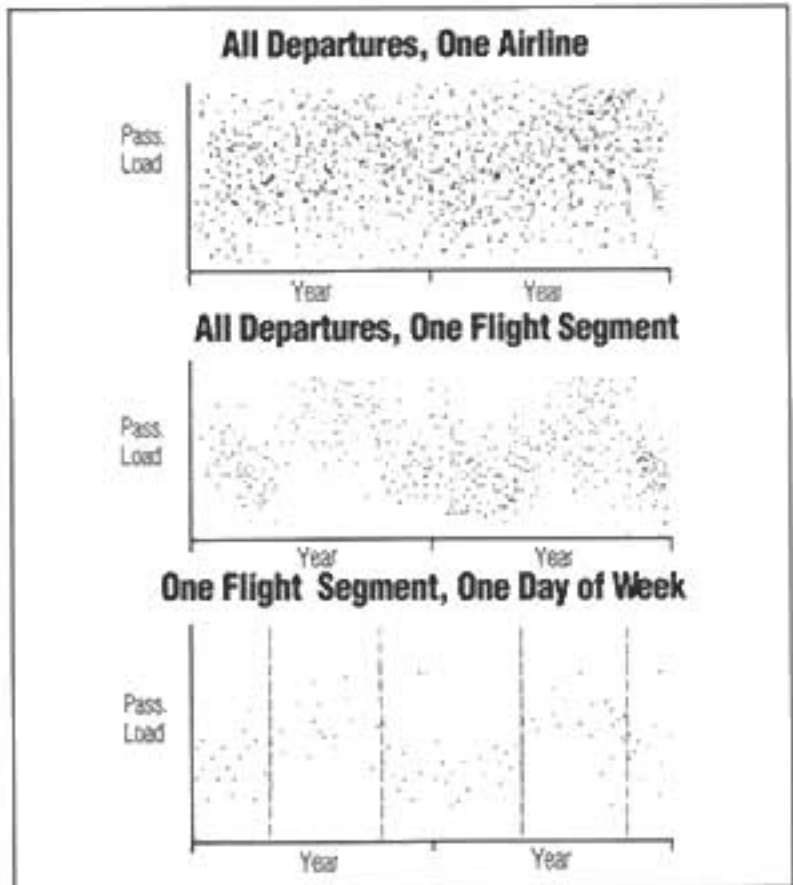


The Problem (cont.)

A tougher challenge for yield management systems is how to deal with inherent or random variation which occurs flight by flight, among flights on the same day of the week, and at the same time of the day.

The extent of variation in demand can be shown by plotting an airline's actual flight loads over time.

Even after the data is disaggregated to one flight segment and one day of the week, scatter still remains from week to week.



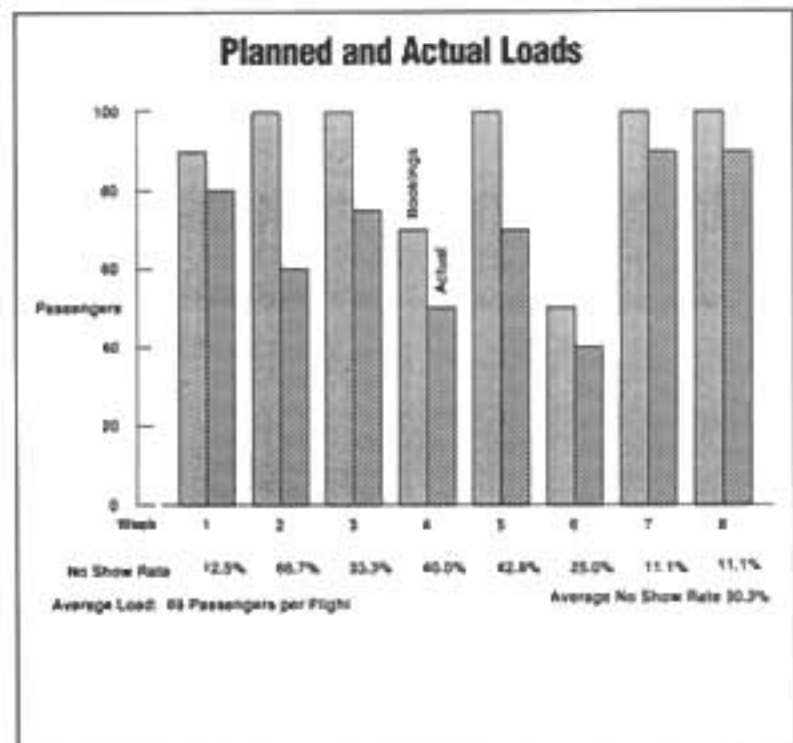
Random variation in demand can be predicted statistically, using modeling techniques in conjunction with the airline's current booking data and recent history for each flight. By looking at the bookings for the previous 8 to 12 weeks of the flight's operation, trends in demand can be determined.

Comparing the bookings with actual loads establishes the rate of "no shows" (people who make a reservation but do not show up for the flight).

Using such data, researchers and system developers are continuing to improve forecasting techniques.

Still, there is an underlying problem for the traditional methods of managing demand:

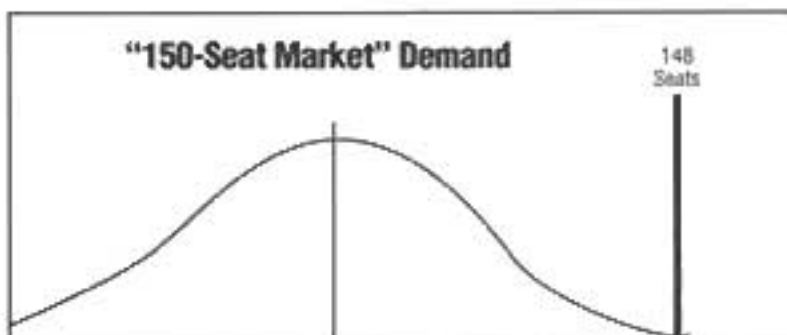
- The supply of seats for a given departure is established at least a month in advance, sometimes up to a year in advance.
- It is nearly impossible for such assignments to reflect the best match of capacities to routes.



The Problem (cont.)


To illustrate the desirability of being able to adjust capacity, consider an example using a normal demand distribution.

In this case, an airplane with 148 seats is deployed in markets where the airplane's average load factor is about 65 percent, and total capacity is rarely exceeded.



About 80 percent of the time, the demands for the 148 seat airplane could be handled by an airplane with 128 seats.

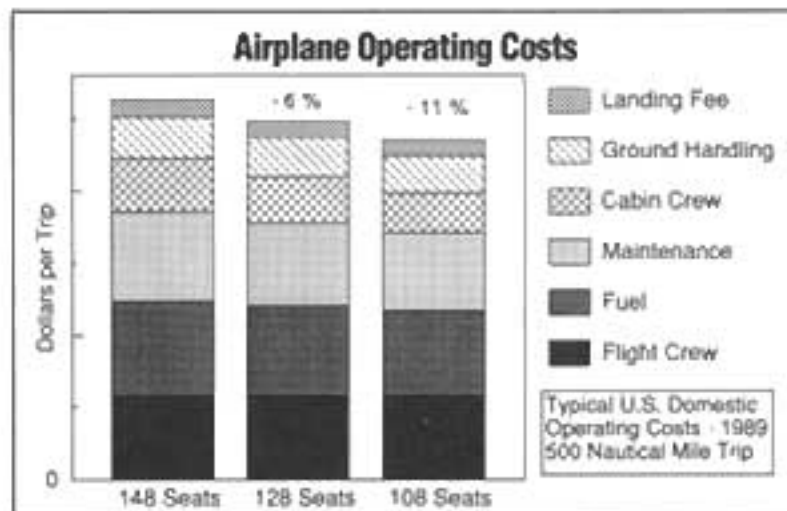
About 60 percent of the time, an airplane with 108 seats could be substituted.



Considering the difference in airplane-related trip costs of the three airplane types, it is easy to understand why an airline would want to pursue the idea of flexible capacity assignment.

This does not imply that the smaller airplanes could be assigned permanently to the schedules of the larger airplane. If that were attempted, spill would be excessive. Demand would start going to competitors, and the airline would no longer have a "150-seat" market.

In summary, it goes back to the quest for the ideal airplane – the mythical **Rubber Airplane**.



The Solution – Demand Scheduling

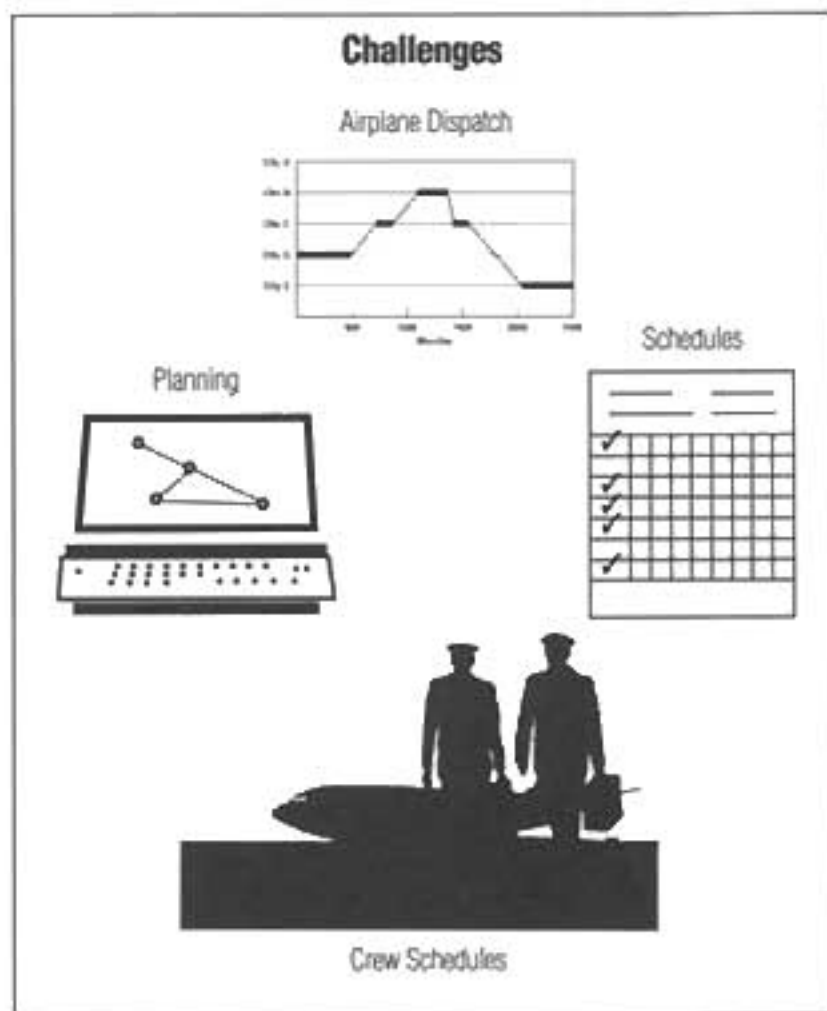
Ironically, it is the existence of truly random variation in demand – the scourge of yield management techniques – which makes the concept of “demand scheduling” so attractive. When demand is higher than expected on one segment, it is likely that demand will be lower than expected on a different segment at the same time.

A few airlines have developed demand-scheduling procedures – shifting airplane capacity between routes to better match demand – but the practice is not widespread in the industry.

The management of capacity by swapping airplanes within a schedule has presented several operational challenges, including:

- Airplane Scheduling and Dispatch
- Ground Handling and Catering
- Maintenance Planning
- Reservation System Issues
- Cabin Crew Assignment
- Flight Crew Scheduling and Utilization

Airlines are generally agreed that the greatest challenge is the latter issue – how to schedule flight crews and achieve an acceptable crew ratio.



The Solution (cont.)

Enter The "Rubber Airplane" (Enabled by D³)

With the advent of common-crew-type-rated airplanes such as the 737 series, ranging in size from 100 to 150 seats, some major problems associated with demand scheduling can be overcome.


For example, a fleet of 737's makes it possible to schedule flight crews independent of airplane schedules. The only requirement is that the airplane scheduled into a particular time slot be one which is common-crew-rated.

Demand Driven Dispatch adds a new tool to the scheduling manager's array of techniques. Airplanes can be assigned to future flights to best match passenger demand for each flight, evaluated in the context of system-wide profitability.

The idea is to deploy the larger airplanes where they can capture high demands, and at the same time promote higher utilization among the smaller airplanes of the fleet, to the extent that they can handle the demand, thus reducing costs.

D³ Benefits:

- Airplane swapping, for load factor improvement
- Utilization shifting, to smaller airplane
- Fleet downsizing, serving the market with fewer total seats



The Solution (cont.)

This concept re-examines the old rule of thumb which held that an increase in passenger load factor had to be accompanied by an increased incidence of passenger spill

Using demand scheduling it is possible to shift to a lower "spill curve", thus *reducing spill* at the same time that *load factor* is increased

Comparing the economic results of an operation using demand scheduling, versus results on the same network and schedule using fixed airplane assignments, an airline will generally see an *increase* in revenue passenger miles and a *decrease* in available seat miles, both contributing to a higher load factor.

The load factor improvement reflects two types of equipment assignment decisions:

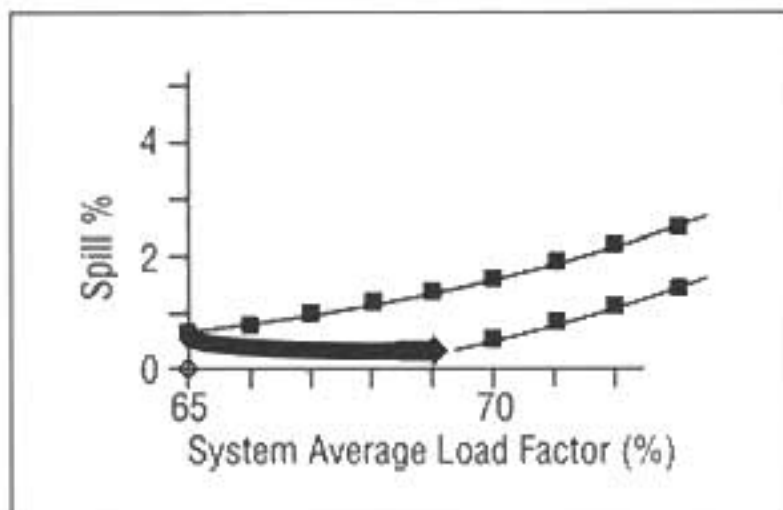
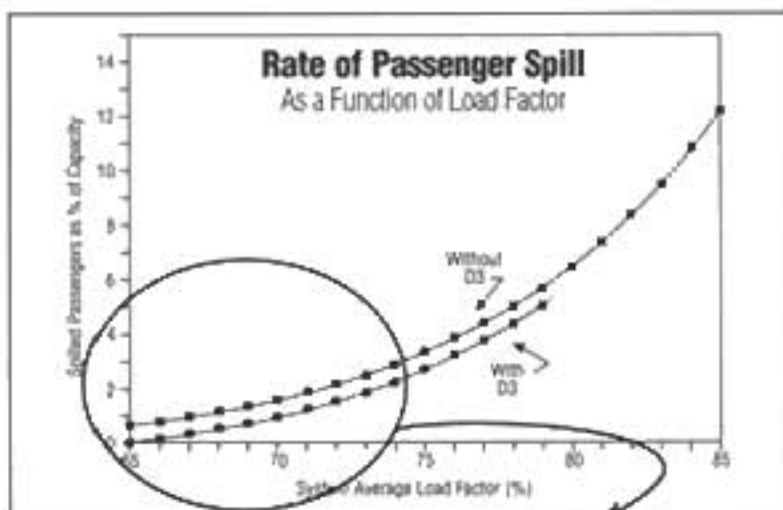
- Statistically high demands can be captured with lower incidence of passenger spill;
- Low or average demands can be accommodated by assigning smaller airplanes to segments of longer distance.

The ability to shift loads to smaller airplanes, on the average, has strong leverage on operating contribution. But this phenomenon is even more powerful when extended to airplane acquisition:

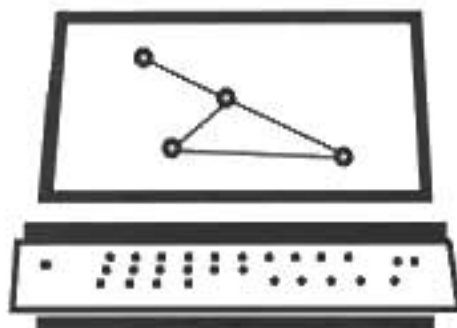
- By planning with demand scheduling in mind, the airline can potentially *downsize the fleet*.
- Future growth in traffic can be accommodated with less growth in total seating capacity.

Beyond these benefits, there are potential savings to the airline from being able to uncouple capacity from the process of developing route networks and system schedules.

Boeing believes that Demand Driven Dispatch is the logical next step for today's yield/revenue/capacity management systems.



Planning Based on D³



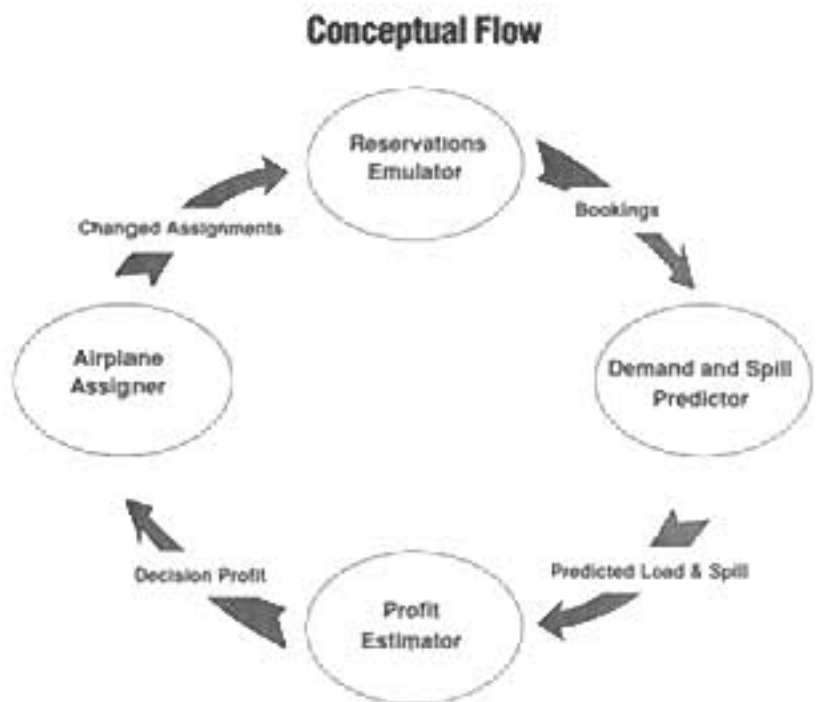
- Future route and schedule development.
- Fleet planning.

Boeing D³ Simulator

Boeing has developed a computer model, the D³ Simulator, to explore the benefits of Demand Driven Dispatch using actual airline fleets and schedules.

The simulator can be described, conceptually, as an interaction between four major functions:

- **Reservations System Emulator**
 - Simulate random passengers
 - Accept bookings and cancellations
 - Turnaway (spill) passengers
- **Demand and Spill Predictor**
 - Observe bookings
 - Forecast demand
 - Predict load and spill for each departure, for each airplane capacity
- **Profit Estimator**
 - Estimate profit (revenue less cost) for each departure, for each capacity
- **Airplane Assigner**
 - Assign airplanes to achieve the best total profit across the network

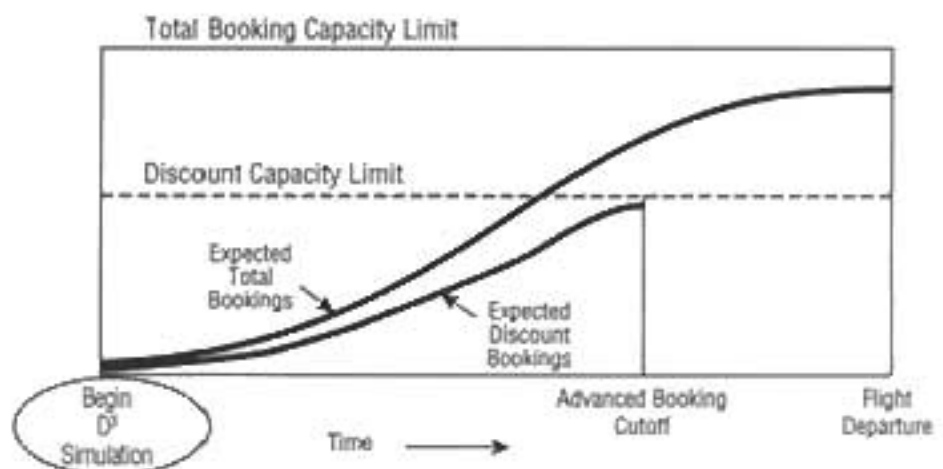


Reservations System Emulator:

The airline's schedules are made available to a simulated computer reservations system. Flights are listed with a generic equipment type (737), without designating a specific airplane type.

Capacity available for sale reflects a tentative assignment of airplane type to each departure, based on estimated average demands.

Simulated Passenger Booking Curve

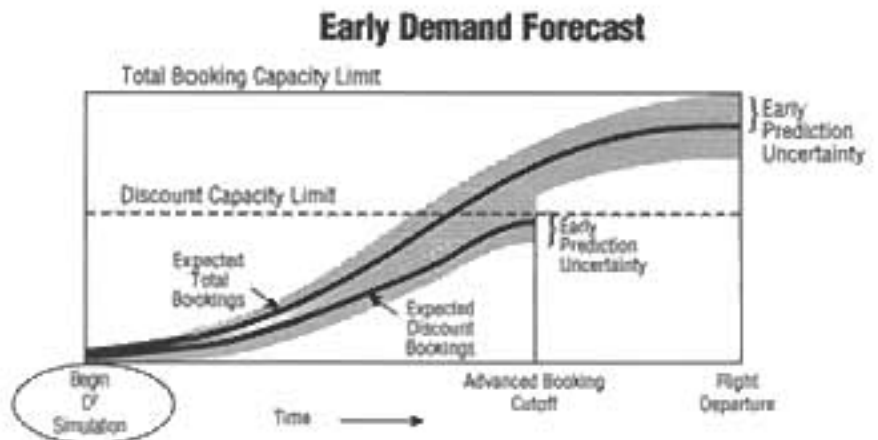


Boeing D³ Simulator (cont.)

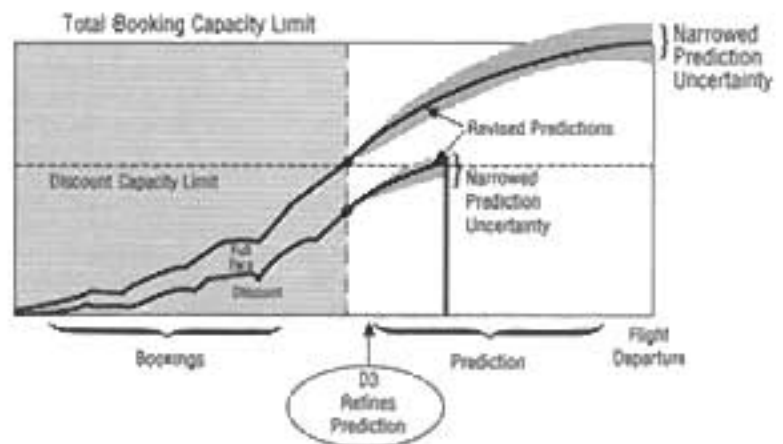
Demand and Spill Predictor:

At some time before the day of service (usually about a month), the D³ model begins daily monitoring of the simulated bookings and prepares a revised prediction of the demand for each departure. This can include independent forecasts for full-fare and discount demand.

If the *predicted demand* for a departure is much higher or much lower than the *previously estimated demand*, that departure becomes a prime candidate for assignment of a different size airplane.



Revised Demand Forecast



Profit Estimator:

Potential contribution to operating profit is calculated for each available airplane type, for each itinerary in the schedule. The profit calculation reflects:

- Estimated revenue from predicted passenger loads (number of passengers times fare or yield)
- Less the variable costs related to the passenger load – including any penalty for spill or denied boarding
- Less the cost of operating the airplane.

Profit Estimator

"Operating Contribution to Profit" Calculation

Estimates:	If flight is assigned:		
	737-500	737-300	737-400
Demand - Spill			
+ Passengers	xxx	xxx	xxx
Passengers x Fare			
- Revenue	\$\$\$	\$\$\$	\$\$\$
- Unit Variable Cost	\$	\$	\$
- Non-Revenue	\$\$	\$\$	\$\$
- Airplane Operating Cost	\$	\$	\$
- Profit	\$	\$	\$

Boeing D³ Simulator (cont.)

Airplane Assigner:

The airplane assignment process determines which airplane type should be assigned to each departure. Input to this process are the decision profits for each combination of airplane type and flight segment. These profits are computed by the profit estimator, based on predictions of loads, spill, and denied boardings.

The assigner seeks to determine an assignment which maximizes total decision profit.

The assigner formulates the problem as a network optimization problem with constraints to ensure that each assignment considered is feasible.

The number of feasible assignments under consideration is determined by the number of "swapping" opportunities available in the network. Swap opportunities arise whenever airplanes of varying capacity are on the ground at an airport at the same time.

While swap opportunities occur throughout the network and throughout the day, particularly rich sets of opportunities occur at connect banks in hub-spoke systems.

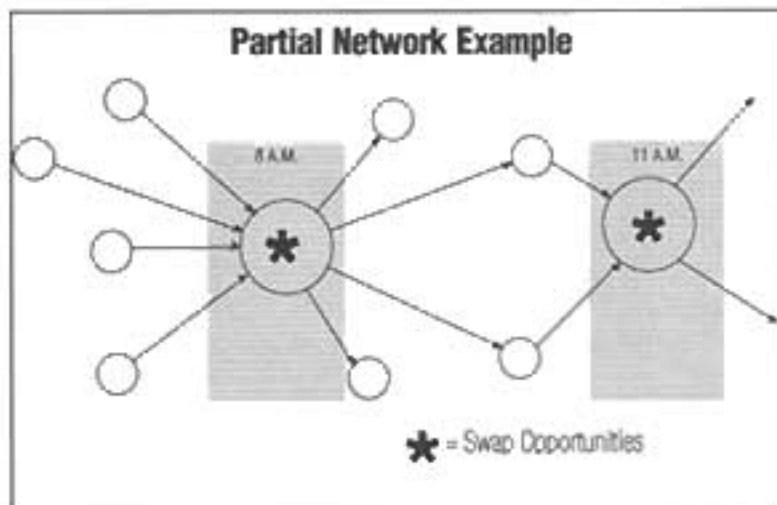
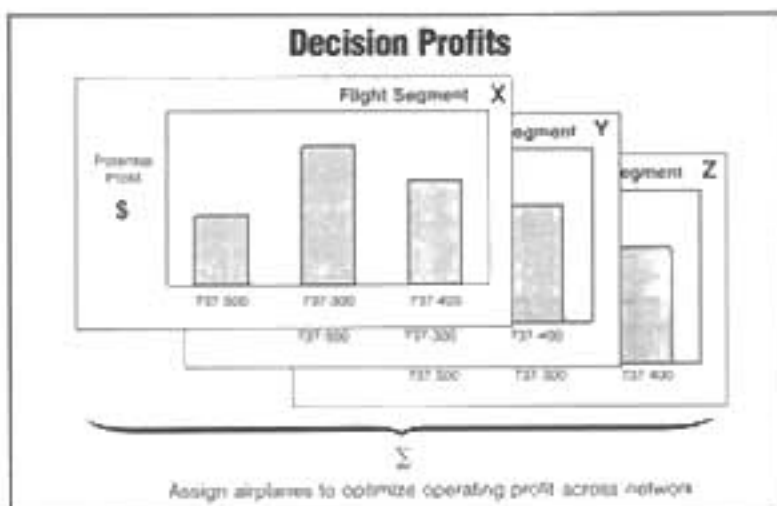
The assignment algorithm, a "heuristic" optimizer, works by iteratively finding the swap opportunities which result in increased total decision profits.

The algorithm was developed to explore the operational practicality of implementing demand scheduling and for use in the D³ simulator to support proof-of-concept studies. It has exceeded our expectations for computational speed, optimality, and extensibility.

As a component of the PC-based D³ simulation, the assigner handles networks consisting of up to 120 cities, 140 airplanes, and 1,000 flights per day. Typical solution times for these networks are measured in seconds.

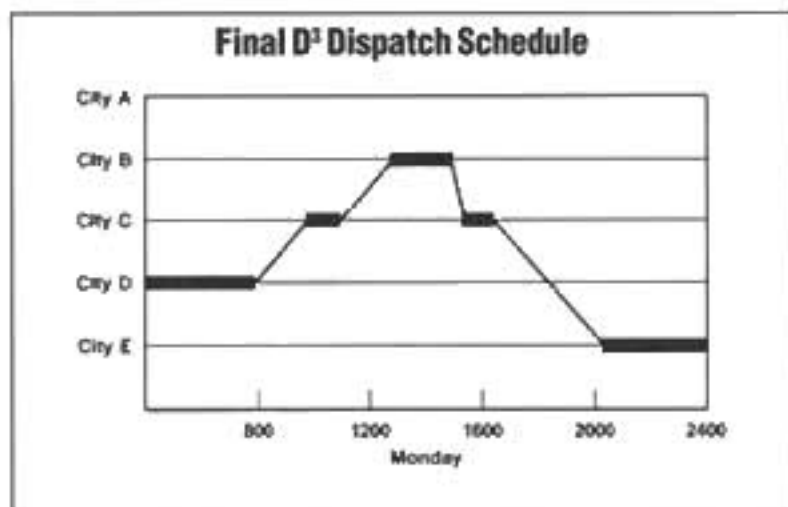
Optimality studies have shown that for problems small enough to be solved formally using linear programming, results generated by the heuristic assigner exceed 99.9 percent of true optimal.

Implementation studies have focused on run-time issues associated with production sized problems and the extensibility of the algorithm to deal with operational constraints. These studies have shown that run times grow slowly as problem size is increased, and that large-scale production problems probably would not require an exceedingly large computer for timely solution. Also, the algorithm has been extended to consider a variety of operational constraints with a small impact on run time.



Boeing D³ Simulator (cont.)

This sequence of activities — observe the bookings, refine the prediction, review/swap the airplane assignments, update the reservations system — would continue until a day or two before the day of service, when the actual dispatch schedule would be prepared and communicated to all stations throughout the system.



Timing of D³ Swap Decisions:

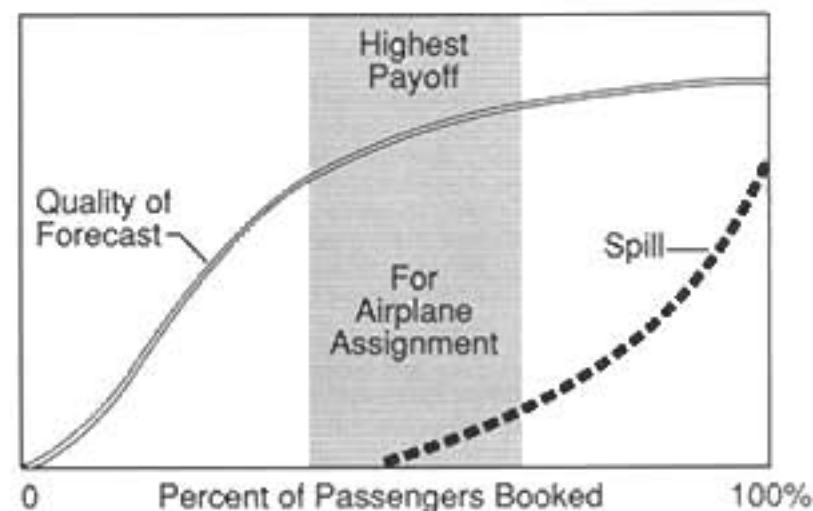
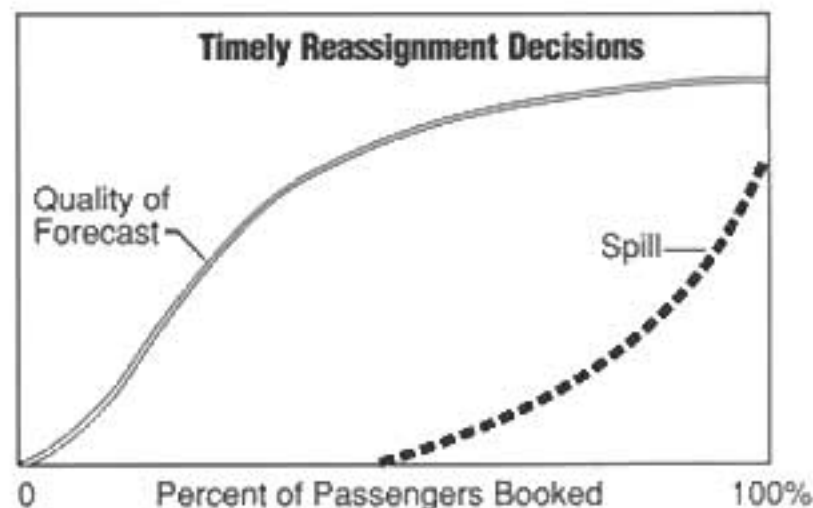
As the flight departure date approaches, passenger demand for each flight segment can be predicted with increasing certainty, giving a better basis for airplane assignment decisions.

The highest payoff for swapping airplanes depends on tradeoff between two objectives:

- It is advantageous to be close enough to the departure date to achieve substantial improvement in forecast quality.
- Capacity swaps must be made early enough in the booking process to avert substantial passenger turnaway (spill).

Typically the best time to swap airplane assignments is during the time period when 40 to 60 percent of passengers have reserved their seats. Such swaps provide high payoff in terms of matching capacity to higher-than-average demand.

Closer to departure date, airplane swaps provide benefit primarily in shifting utilization from larger to smaller airplanes, when either airplane can handle the forecast loads.



The Payoff

Demand Driven Dispatch Financial Benefits

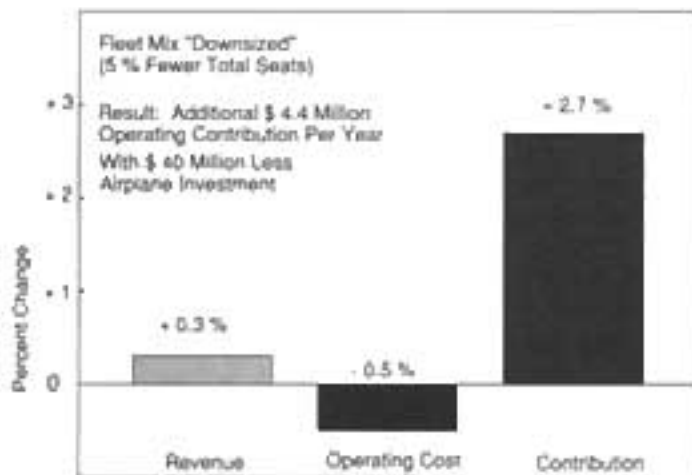
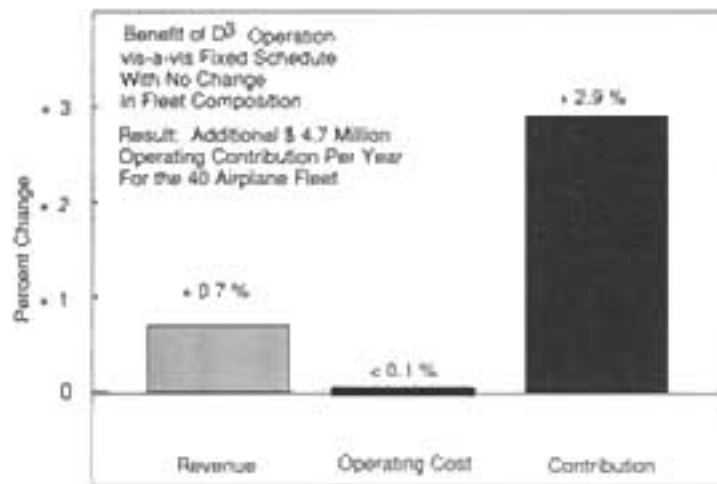
To illustrate the financial benefits of D³, a case study example shows the potential for improving contribution by use of demand scheduling in a 40-airplane system consisting of 22 cities and 244 flights per day. A description of this case study is included on pages 15 and 16.

The case study results, shown graphically, illustrate the impact of the D³ concept.

Based on several case studies with airlines, as well as simulated exercises, the financial benefits of Demand Driven Dispatch can be expected in the ranges shown in the box, bottom right.

Several studies have resulted in even higher net dollar benefits per airplane.

Case Study Financial Results



Beyond the Simulator:

In an operational application, the airline's own reservations system would replace the D³ Reservations Emulator.

The function of refining demand and load predictions would need to be incorporated into the airline's yield management system (if not already available in the system).

The functions of the D³ Profit Estimator and Airplane Assigner would be adapted to the airline's operations, to work in conjunction with reservations, yield management, dispatch, maintenance scheduling, and other systems.

Area of Improvement

Typical D³ Payoff

Demand Matching

1 to 3% increase in load factor.

Utilization Shifting

1 - 2 hours per day
shifted to smaller airplanes.

Fleet Downsizing

5% to 15% fewer seats

Operating Contribution

Additional \$50,000 to \$200,000 per
year, per airplane, contribution to
operating profit.

D³ in Operation

For many airlines, the elements required to implement a Demand Driven Dispatch system are already in place. They include:

An airplane family with common crew rating

Although it is possible to accrue benefits without a common-crew-rated fleet, the availability of a common rated fleet significantly reduces crew cost, training, and administrative complexity.

Computer Reservation System (CRS)

The ability to forecast demand and observe booking trends, by class of service, is critical to the concept of shifting capacity on an as-needed basis. These activities can be done off-line or during off hours so as not to impact the real-time operation of the computer reservation system.

Analytical Computing Power

The airline has to solve a giant assignment problem on a daily basis. Although we run the D³ assignment algorithm on a large personal computer system at Boeing for case studies of impressive size, the computing power needed by the airline would depend upon the size of the network and schedule.



D³ Case Study

Shown here is an example of how D³ can work in an airline environment. The schedule shows one time period for which five airplanes are scheduled to depart from city "A".


The airplanes used are

Airplane	Quantity	Capacity (Dual Class)
737-500	2	108 seats
737-300	2	128 seats
737-400	1	148 seats


The assignment of airplanes is based on historical average demand

By using the D³ model, together with a refined demand forecast, new airplane assignments have been recommended for all but one segment. The new assignments are based on increasing the operating contribution across the network

Planned Schedule
Historical Average Demand (μ)



D³ Schedule
Forecast Actual Demand (D)



Revised Schedule Based on D³ Model

Segment	Original Airplane	D ³ Assigned Airplane	Reason for Change
A - B	737-300	737-500	Demand lower than expected. Smaller airplane assigned to reduce cost.
A - C	737-300	737-400	Demand higher than expected. Avoid spill. Also, dispatch the largest airplane on shortest segment to reduce cost.
A - D	737-500	737-300	Demand lower than expected. However, smaller airplane (737-500) was better utilized on longer flight from A to F.
A - E	737-400	737-300	Demand can be accommodated by smaller airplane. Reduce cost.
A - F	737-500	737-500	No change.

D³ Case Study (cont.)

In this example, the D³ model adjusted the scheduled airplanes to increase revenue and reduce operating cost. It accomplished this by assigning airplanes to avoid passenger spill and to shift utilization to the smaller airplanes when they could handle the forecast loads.

Demand Driven Dispatch Financial Results (Per Day)			
Fleet	Published Schedule	D³ Schedule Same Fleet	D³ Schedule Adjusted Fleet
737-500	12	12	20
737-300	14	14	10
737-400	14	14	10
Passenger Load Factor	62.6%	64.6%	66.2%
Spill (Passengers)	1,026	893	1,031
Denied Boardings	26	26	33
Revenue (\$000)	\$1,943	\$1,957	\$1,945
Operating Cost (\$000)	\$1,492	\$1,493	\$1,482
Operating Contribution (\$000)	\$451	\$464	\$463

Airplane Investment (\$ Millions)	\$1,172	\$1,172	\$1,132

The network example shown on the preceding page was but one snapshot out of the thousands of assignment decision points included in the case study.

The result of the study, using the D³ model, was an improvement in load factor and revenue, with no increase in spill. The D³ operation, with no change in fleet, resulted in 2.9% improvement in operating contribution.

As an added benefit, there is the opportunity to reduce the average airplane size within the fleet, thus reducing investment. Comparing the third column results against the first column, there is a 5% reduction in total seats, a \$40 million up-front savings in airplane investment, and almost the same operating improvement as obtained in column 2.

These case study results are comparable to a number of studies done for airlines.

If you would like to know more about the D³ concept and the specific benefits for your airline, please contact us:

Ron Caldron	(206) 237-6971
Don Kitch	(206) 234-2730
Virginia Sweetland	(206) 234-4318

Airline Market Analysis
Mail Stop 75-14
Boeing Commercial Airplanes
Seattle, Washington 98124

Or ask your Boeing Sales Director to arrange
a discussion

