## Chapter 1

# Introduction

We live in a complex and dynamic world where uncertainty fittingly characterizes its intrinsic nature. In this world, change is constant and everything else is variable.

What is the business environment in this rapidly transforming world? Advanced information technology through global networks has significantly influenced today's business by increasing competition, eliminating boundaries, and improving access to information. The marketplace in the 21st century is technology-driven and fiercely competitive. To succeed, or even just survive in this new era, it is essential to maintain a commitment to effective and scientific management of all available resources. Since our business environment is evolving at a rapid pace, our decisions should also adapt to the change.

Traditionally, people emphasize planning – making a detailed and complete blueprint for actions to gain the highest value. Needless to say that planning is very important. Making a sound plan before taking actions is a fundamental principal that guides peoples' practice in various disciplines. However, a good plan is only half of the process. No matter how superior a plan is, in the execution phase, various unanticipated events will disrupt the system and make the plan deviate from its intended course and even make it infeasible. How to cope with disruptions? How to reach our goals while minimizing all the negative impact caused by disruptions? How to get back on track in a timely manner while effectively using our available resources? These are the essential topics investigated in the field of "disruption management."

Applications of disruption management range from airlines and manufacturing systems, to telecommunications and educational in-

DISRUPTION MANAGEMENT - Framework, Models and Applications

stitutions. The value of disruption management has been demonstrated via multiple sources, and has recently been recognized by the OR/MS community. For example, the INFORMS 2002 Franz Edelman Award was given to the work undertaken by Continental Airlines to deal with the problems of crew disruptions (Yu *et al.* 2003). The 2003 IIE Outstanding Publication Award and 2002 IIE Transactions Best Paper Award were given to research conducted by Bard, Yu and Argüello (2001) on reconstructing aircraft routings in response to aircraft groundings delays.

Disruption management is dramatically exemplified by the event that a lightning in March 2000 caused a ten-minute fire at Philips Electronics manufacturing plant in New Mexico, which left the supplier short of a critical chip for several weeks. The two telecommunication companies, Nokia and Ericsson, which directly source from Philips, took different remedial actions. The outcomes were drastically different: Nokia gained 3% of the handset market mostly from Ericsson while Ericsson lost $1.68 billion, largely caused by the event (The Wall Street Journal, 1/29/2001). The other example is the 2003 U.S.-Canada blackout – a massive power outage that occurred in parts of the northeastern United States and eastern Canada on August 14, 2003. It was the largest blackout in North American history, affecting an estimated 10 million people in Southern Ontario, Canada and 40 million people in eight U.S. states. This event severely affected the logistics of the companies residing in the affected area as well as their customers. For example, to reduce loss and ensure operational continuity, a Chinese publishing company in Beijing quickly rescheduled its production and switched its orders on high-quality paper supplies from a US company in New Jersey to a company in Spain.

The following case study gives a thorough coverage of how disruption management concept is applied and what impact it creates.

## 1.1 Case Study: Disruption Management at Continental Airlines

There is not a single day at any time of the year that airlines are able to operate smoothly and completely as planned. The buzz word "irregular operations" in airlines' terminology usually refers to disrupted operations from the originally planned schedule. The disruptions come from all possible causes, ranging from inclement weather, aircraft mechanical problems, crew sickness, and union strikes, to the now more frequently faced security-related incidents. It might be more proper to define "irregular operations" as regular operations since there has never been a "regular operation" in any day of any airline's history.

Moreover, due to severe competition, airlines' profit margins are sharpened razor thin, with large sums of fixed costs spent on airplanes, fuel, and facilities. In good times, when airlines have stable RASM, or Revenue per Available Seat Miles, they make money in regular operations. But in irregular operations, airlines lose money. So the real trick to making money is to get back to regular operations as quickly as possible, and minimize the loss and negative impact in the recovery process. This is the "disruption management" concept accepted in the airline industry.

### 1.1.1 *The Storm of the Century woke Continental Airlines up*

The Storm of the Century in 1993 became the catalyst for Continental Airlines to search for new strategies to deal with disruptions. In March 1993, the super blizzard hit the United States, the worst since the legendary blizzard of 1988, claiming 240 lives, affecting 26 states, and causing damages of approximately $1 billion. Twenty inches of snow were dumped in the Southeast, 11 tornadoes were spawned in Florida alone, and hurricane-force winds of over 75 miles/hour resulted in grounded aircraft up and down the eastern seaboard for days. One of Continental's hubs – Newark Airport – was closed for almost two days.

It took Continental five days to dig out from under the storm. Employees located airplanes by brushing the snow off the planes' identification numbers. Crew managers found crews by calling the airports to find out where they had been sent for accommodations. Some crews stayed together and others were dispersed among two or three different hotels. It took days for Continental to figure out the location of all its crews. Most flight crews tried to call in to the operations center but found the phone lines jammed. From an operational standpoint, Continental completely lost control of its operations.

Such a disaster woke Continental up. Continental's management began to reexamine its operational processes. The senior management pulled thirteen employees from their duties in the operations center and formed a taskforce for improving recovery operations. Anna White, director of Crew Technology, set out to find a viable solution to quickly and efficiently get its flight crews in place to fly following a major disruption to operations.

### 1.1.2 *CALEB Technologies responded to the call*

Recognizing the immediate need, Continental wanted to buy an off the shelf product to handle major crew disruptions, customize it, and have it deployed in six months. However, it did not quite work out that way. The taskforce visited both domestic and international airlines, to see what software others were using to handle the problem. It turned out that other airlines were looking as well. Then, Continental talked to the major software vendors in the airline industry and found that all the existing software was built on prevailing pairing optimization techniques – suitable for planning optimization but not for real-time disruption management. Dr. Tom Cook, the former President of Sabre Decision Technologies and INFORMS president, acknowledged that people have tried and failed many times to develop such a system, and no similar system existed in the market place at that time. It was not because of lack of interest, but due to the sheer complexity of dealing with a large-scale complex problem, sophisticated government regulations, and complicated disruption scenarios.

Then, through its information technology provider – Electronic Data Systems (EDS) – Continental discovered Dr. Gang Yu, and his newly founded CALEB Technologies. As a professor at the McCombs School of Business at The University of Texas at Austin with strong expertise in operations research, Dr. Yu also had ample airline experience. He had traveled to United Airlines' headquarters in Chicago numerous times and sat in on the System Operations Center. For countless numbers of hours, he observed and investigated how the operations managers handled disruptions. As a result, he helped United Airlines build the airline industry's first disruption management system – the Delay and Swap Advisor in 1991. As Chairman and CEO of CALEB Technologies, Dr. Yu responded to the call and was determined to build a real-time decision support system – CrewSolver – to manage crew recovery under disruptions.

Larry Kellner, President of Continental Airlines, recollected, "We would like to do projects that make good business sense and benefit our employees. Matching crews and aircraft to meet the schedule is always a complex problem. We couldn't find a tool that fit our philosophy and approach, so we partnered with CALEB to build one. We are willing to try innovative techniques and take risks when we see the potential. We saw a huge potential here." With a vision to lead the airline industry, Continental invested in Dr. Yu and his then three-person company. To mitigate risk, Continental also bought an insurance policy for Dr. Yu to protect the potential multi-million dollar losses if Dr. Yu could no longer lead the project due to unforeseen contingency.

### 1.1.3 *Background of Continental Airlines*

Continental Airlines, a major United States air carrier, transports passengers, cargo, and mail. It is the fifth largest United States airline and, together with its wholly owned subsidiaries, Continental Express and Continental Micronesia, operates more than 2,000 daily departures to 123 domestic and 93 foreign destinations.

Continental operates its domestic route system primarily through its hubs in the New York metropolitan area at Newark International Airport (abbreviated as EWR), in Houston, Texas at George Bush Intercontinental Airport (abbreviated as IAH), and in Cleveland, Ohio at Hopkins International Airport (abbreviated as CLE). This hub system allows it to provide passenger services between a large number of destinations more frequently than it would by servicing each route directly. Such a system also allows Continental to add service to a new destination from a number of cities, using a limited number of aircraft. Each domestic hub is in a well-populated area and large business center, ensuring a high volume of passenger traffic. Continental serves more non-US cities than any other US carrier, including cities throughout the Americas, Europe, and Asia. It has more than 50,000 employees, including 4,000 pilots and 8,000 flight attendants.

Continental's system operations control center (SOCC) is located at its headquarters in Houston, Texas. At the SOCC, Continental personnel monitor operations, track the execution of schedules, anticipate disruptions, and determine the recovery from disruptions. The SOCC provides a central location for making all decisions affecting airline operations, including customer service, crew scheduling, aircraft routing, maintenance scheduling, and dispatch. When disruptions occur, SOCC personnel change the flight schedule, perhaps canceling or delaying flights, route aircraft to support those changes, and finally reassign crew to fly the new schedule. Although they make these decisions sequentially, they do not make them in isolation. They use advanced systems to view the impact one decision may have on another. The operations managers who change the flight schedule and route the aircraft consider the impact on passengers, crew, and required scheduled maintenance in making these decisions. They confer with customer service representatives, crew coordinators, and maintenance routers when making recovery decisions. After the operations managers determine the new flight schedule and aircraft routings, the crew coordinators take over to assign crew to uncovered flights and recover crew back onto their original schedules.

### 1.1.4 *The Go-Forward plan and the CrewSolver project*

Seven years ago, Continental was coming out of bankruptcy. Continental's Chairman and CEO Gordon Bethune actually had to call Boeing and say, remember those deposits we gave you for new airplanes? We need them back to make payroll. That's where Continental was. To get the airline out of bankruptcy, Bethune came up with a Go Forward plan. This plan has four cornerstones: Fund the Future, Make Reliability a Reality, Fly to Win, and Work Together.

Does the CrewSolver fit the Go Forward Plan? Janet Wejman, Continental's Chief Information Officer, and Anna White examined where Continental was from a crew technology standpoint to make sure that CrewSolver was in sync with this Go Forward plan. It was easy to see that Continental needed an operations research tool to Make Reliability a Reality. At that time, Continental was considered the worst airline. As illustrated by Bethune's book, *From Worst to First*, it was quite obvious that Continental was the worst airline so much so that the second worst airline looked good compared to it. What Continental needed was operational consistency, in which crews play a major role. So Wejman approached Bethune, and told him that Continental needed the CrewSolver tool to help them get there, and explained how and why. And even during those bankrupt times, Bethune went for it because he understood the value of making operations run more smoothly.

The next cornerstone of the Go Forward plan is Fund the Future. Continental's management knew that if Continental's operations ran more smoothly, they could put dollars back on the bottom line. Such benefits can be seen in the reduction in hotel stay for crews, crew overtime pay, customer service agents, and reservations agents; in flying the expensive leased aircraft to generate revenue, versus sitting on the ground; and in not canceling flights and potentially losing the high yield revenue passengers, who pay full fares to fly. That is where CrewSolver has enabled Continental to put dollars back to the bottom line, and helped Fund the Future.

The next cornerstone of the Go Forward plan is Fly to Win. Continental's goal is to be in the top 25% of major airlines for reported profit margin. Continental's management knew that CrewSolver could contribute, by helping them maximize completion factor. By successfully crewing all operating flights after a schedule disruption, no additional flight cancellations will be incurred as a result of crew coverage. That, in turn, minimizes lost revenue from additional cancellations.

Working Together is a big component of Continental's Go Forward plan. Continental's Golden Rule is to treat both the internal and external customer, with dignity and respect. Continental believes that both their paying customers and their crew members deserve open and honest communication. CrewSolver allows them to take an irregular crew situation, quickly assess their options, choose the best course of action, and implement that plan in a matter of minutes. This means they can provide consistent and reliable crew information internally, which translates into better communication to both their crews and their customers.

Thus, CrewSolver addresses all the major points that are key success factors in the Go Forward Plan. However, can the system bring the needed return on investment (ROI)? Dr. Yu and his team spent two months building a prototype system and used Continental's MD 80 fleet as the test bed. By comparing the decisions from Continental's crew coordinators and the optimal decisions generated by the prototype system on some typical disruption scenarios, they showed that the ROI should be within eight months. Equipped with all the above, Continental's top management gave the go-ahead.

### 1.1.5 *Challenges in modeling, solutions, and real-time technology*

Now the challenge has shifted to Dr. Yu and his CALEB Technologies to build a real-time decision support system to face the complex and dynamic environment, take all the components into account, and handle all the possible scenarios. The Delay and Swap Advisor Dr. Yu helped build for United can only handle aircraft and not crews. The crew recovery problem is several orders of magnitude more complex due to the significantly higher number of crew members than aircraft and a larger quantity of crew restrictions due to safety considerations.

The detailed modeling and solution methodologies of the problem will be discussed in later chapter. To put it simply, the crew recovery problem can be described in the following manner: when faced with disruptions, how do airlines get the crew back to planned assignments in a timely manner? The goal of the crew recovery problem is to minimize the recovery cost to cover all flights remaining in the schedule with the required and qualified crew, while keeping the crew on their original assignments as much as possible. The crew recovery problem is proven to be NP-hard, a term to describe the intractability of the problem due to its exponential growth in required computational time when the problem size increases. Usually, the number of variables is huge due to the combinatorial number of possible pairings each crew member can take. In Continental's situation, there are millions of discrete variables.

Dr. Yu formed a team consisting of operations research analysts and information technology specialists to tackle the problem. They had breakthroughs in modeling and solution techniques, and multiple patents were generated from their work. Some of their research findings are included in this book.

The next challenge is acquiring massive and accurate information in real time. The CALEB and EDS team designed the system such that the CrewSolver optimization server contains an in-memory data store that represents the current airline operations. The optimization model is always maintained up-to-date in the data store. This way, when a request is received, all processing is dedicated to solution generation, and no time is wasted on external database transactions. It runs continuously, 24/7. Upon receiving a request from the crew graphical user interface, the server launches a new process to set up a problem scenario based on the data input by the user and the data store. Real-time data update is done through messaging. Optimization model update is triggered whenever data is updated.

From a business perspective, the following is the process CrewSolver goes through to solve a disruption problem. CrewSolver takes up to the minute operational data, looks at real time disruptions, applies current crew constraints, and develops optimal and legal crew solutions. It then provides those details to the users, quickly implements and updates all crew records, and pushes that information to other operational systems, which in turn updates the reservations centers and airports. It also pushes the crew changes to the web, so that crew members can review and acknowledge changes to their schedule. Using CrewSolver, this process takes anywhere from 3 to 30 minutes, depending on the recovery window and the size of the problem. What is amazing here is that if the same situation were handled manually, it would take anywhere from an hour to a day, and it would not be an optimal solution. Continental also uses CrewSolver for what-if scenarios, so as the SOCC develops flight delay and cancellation strategies, they can take those proposed changes and determine the crew impact. That is incredible and something Continental could not do in the past. The problem was just too large and the possibilities far too time consuming.

DISRUPTION MANAGEMENT - Framework, Models and Applications

With all the optimization and information technology developments, CrewSolver is put to the ultimate test. Will it work? And will Continental get the value back?

### 1.1.6 *The payback and the testimonials*

While CrewSolver can be used for all types of crew disruptions, there were a few events in 2001, which highlight the variety of disruptions handled by CrewSolver.

The first event is New Year's Eve 2000, the East Coast's worst storm in five years. Since Newark is one of Continental's hubs and a major focus of its operations, Continental looked at the weather projections, and that Friday night reduced Saturday's schedule by 35%. Unfortunately, the weather was much more unforgiving, and the storm dumped 18" of snow overnight. This completely closed the Newark airport for a number of hours. Continental regrouped and cancelled an additional 30% of the EWR operation that morning. They were able to very quickly put a plan into place, implement it, and advise the crew members of their changes. The point here is that Continental recovered faster than its competitors, and had confidence that they could fly the reduced flight schedule. They got the passengers where they needed to go as quickly as possible and, in the process, picked up other passengers because Continental was flying again and in record time.

The next is the "reluctant storm" in March 2001. It went from being the East Coast Storm of the Century to a massive storm for the area, to a moderate storm impacting certain areas, and finally to a moderate snowfall. It still had a large impact on the system, but nowhere near the original forecasts. The point of mentioning this storm is to show how Continental reacted. In the past, given this forecast, Continental would have begun canceling flights early, given the magnitude of the crew updates. With CrewSolver, Continental delayed canceling flights until they had a good picture of the actual weather impact, and eventually cancelled only 141 flights. CrewSolver allowed Continental to develop the best information possible, before determining what to do. This was a major win-win situation, since they were able to minimize irregular operations from a series of days to only one day.

Tropical Storm Allison in June 2001 was another big opportunity for CrewSolver. The storm hit Houston for three days, poured over 30 inches of rain, and resulted in Houston being declared a national disaster area. It was not good for Continental, since Houston is one of their three domestic flight hubs and accounts for 30% of their total flight operations. Plus it is a major crew hub, so this storm had the potential to cripple Continental's entire operation due to lack of crews. Unfortunately, in this case there was no master recovery plan. Continental just had to take it a day at a time as the flooding continued. Continental was able to maximize the use of flight crews in EWR and CLE, and minimize the impact of stranded Houston crews on flights operating outside of Houston. That was another big win for Continental.

The most important test of the CrewSolver system's abilities came on and after Tuesday, September 11, 2001, when the FAA closed the airspace over the United States and diverted all planes to the nearest airport following the attacks by terrorists using four aircraft from major US carriers. Suffice it to say that, from a crew perspective, airlines had never before faced a disruption of this magnitude. The first significant challenge was to bring the airline back up after a 3-day period shut down. In one CrewSolver solution alone, Continental repaired over 1,600 pairings on the 737 fleet, which was by far our largest solution ever. Then Continental reduced its flight schedule for the balance of September by 20%, and they needed to reschedule all flight crews for a 14-day period. That is over 13,000 crew events. Although CrewSolver was designed to handle a 3-day operational window, CALEB was able to quickly modify the software to look at this 14-day extended window. Fortunately, when CALEB built CrewSolver, they did it right. When it was finally finished, a total of 14 days of disruption was repaired, and basically all crews for the balance of September were re-planned. Without CrewSolver, Continental would have spent the better part of those two weeks manually recovering the operation one day at a time, which would have resulted in a much larger percentage of overall delays and cancellations.

Here is how Continental fared as compared to other airlines, in terms of start-up delays for those first few days. Continental did cancel 20% of its operations, but they were confident that they could fly the remaining 80% with their aircraft and crews, and do it from day one. Figure 1.1 shows Continental's performance. The delays encountered were due to security, and not operational reasons. This helped to minimize the revenue losses, as well as reduce crew costs. Other airlines were not as fortunate, and incurred 30 to 50% delays in the first few days of flight schedule start-up.

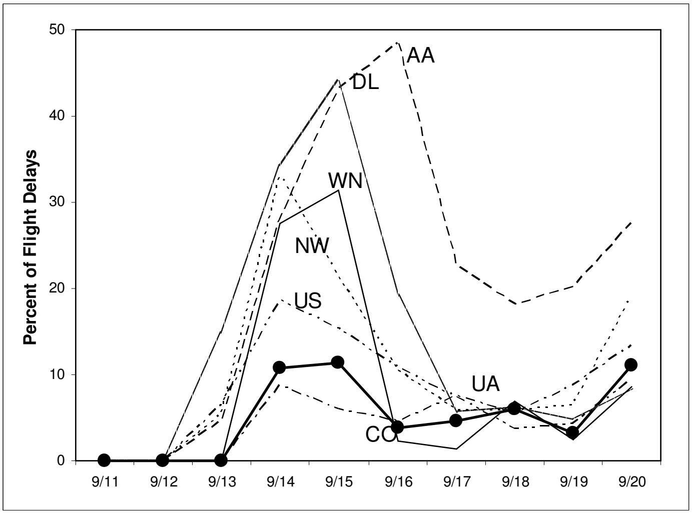

Fig. 1.1 Comparison between Continental and other major airlines after 911 Event

As these events evidence, applying operations research (OR) to the crew recovery problem at Continental, and integrating CrewSolver into daily operations, has resulted in a number of qualitative and quantitative benefits. The OR methods used in CrewSolver allow one to search through millions of possibilities for the best overall crew solution in a matter of minutes. We cannot emphasize enough the time criticality associated with irregular operations, or the positive impact that a good, quick crew plan can have on the entire recovery process. Every hour is money to an airline.

In looking at the 2001 statistics, it can be concluded that using CrewSolver saved Continental anywhere from 1 to 5 million dollars per major disruption. This does not take into account the dollars saved during daily minor disruptions.

Just looking at the four major events from 2001, CrewSolver benefited Continental Airlines to the tune of approximately $40 million dollars. Far and away, the biggest benefit was during September, where Continental sustained major irregular operations for 18 days. To put that $40 million into perspective, in 2000 Continental's net revenues were $341 million, while in 2001 Continental's net losses were $95 million.

Figure 1.2 illustrates Continental's load factor premium compared to the rest of the industry. In part, this is due to Continental's operational consistency. CrewSolver has a significant impact on that consistency.

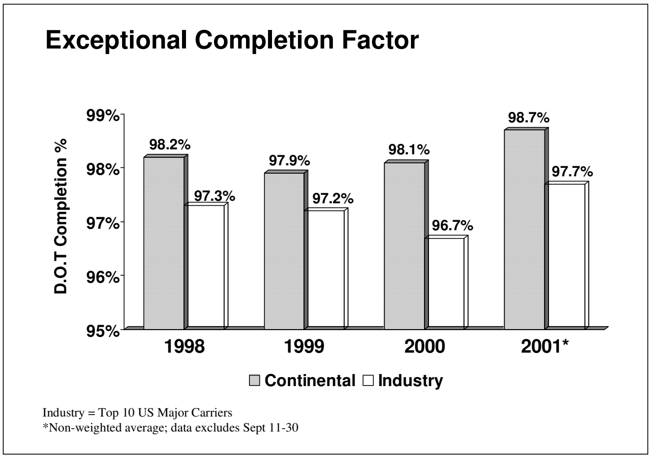

Fig. 1.2 Comparison on the completion factor between Continental and other major airlines

Figure 1.3 shows Continental's Department of Transportation (D.O.T.) arrival statistics. One can see that Continental is an industry leader in on-time performance and it keeps improving over the years especially after adopting CrewSolver. That again makes a difference to the bottom line. So being an on-time airline, and completing the flights promised to fly, makes a large difference in profitability.

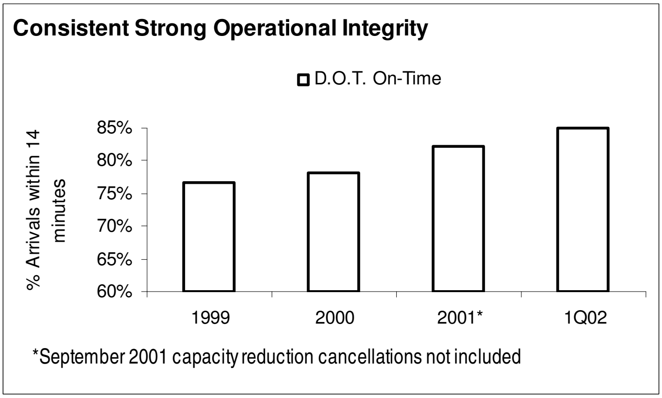

Fig. 1.3 Continental's Department of Transportation on-time performance

Big numbers are all around, which makes real time decision support systems a very valuable tool, not only at Continental, but in the airline industry as a whole. Following Continental's great success, several other airlines have also adopted CALEB recovery systems. Southwest Airlines, Northwest Airlines, and JetBlue had their customized implementation of CrewSolver in production. Delta Air Lines is in the process of installing CrewSolver.

### 1.1.7 *The vision, the Franz Edelman Award, and the bright future*

On May 6, 2002, Wejman, White, and Dr. Yu presented Continental's remarkable "recovery" story – "A New Era for Crew Recovery at Continental Airlines" – before a panel of judges at the INFORMS Practice Meeting in Montreal. As one of six finalists for the annual Franz Edelman Award for Achievement in Operations Research and the Management Sciences, the Continental and CALEB team was competing in the "Super Bowl" of OR/MS against a worthy field that ranged from French automaker PSA Peugeot Citroen to U.S. candy maker Mars, Inc. A day later, the team presented an encore performance, this time for the world to see as the 2002 winners of the coveted Edelman Award.

"When we started all of these seven years ago, we had a vision," White said. "We thought we could make a contribution to the company, but I had no idea it would be such a major contribution – and so timely. It's especially rewarding winning the Edelman because it demonstrates that a group outside the airline industry has recognized the fact that this product brings real value to Continental."

Wejman reflected upper management's enthusiasm for operations research solutions to complex operational problems at Continental, an enthusiasm that extends all the way up to Chairman and CEO Gordon Bethune and President Larry Kellner. "The airline industry has long understood the value and impact of operations research, and we're proud to be leading the way," says Wejman, noting that Continental was the first airline to adopt the cutting-edge recovery tools. "Larry Kellner and Gordon Bethune understand the value of technology at the airline, and they have been great supporters of our efforts throughout my tenure." She further affirmed, "As CIO, I can tell you that no other single application in my tenure at Continental has saved the company more money than CrewSolver."

Dr. Yu added, "Operations research is at the core of CrewSolver; it's what creates real value. If you just build an information system with databases and a pretty interface, you can only achieve a fraction of what you can do with operations research. With OR, we can really help airlines manage their manpower and save tens of millions of dollars. That's the value of OR. That's the future of OR."

Bethune and Kellner both appeared via videotape providing testimonials. Bethune told the judges that CrewSolver gives Continental "a competitive edge" by helping the airline get the right crews with the right planes back in the air during abnormal operations. "It optimizes not only our performance, but it makes sure we provide the highest levels of customer satisfaction. Technology is wonderful when it gives you an edge. CrewSolver really hit a home run."

Added Kellner: "From a technology perspective, CrewSolver was a key advancement for our company. By applying advanced optimization techniques to real-time business problems, we achieved our objective. The flexibility of the tool could not have been tested more than it was by the event of September 11th. Yet, we had a new schedule in place by October 1 and CrewSolver passed the test with flying colors. From a financial perspective, the project was a homerun. It reduces irregular operations recovery costs from a crew perspective and it reduces passenger re-accommodation costs, and it also reduces cancellations and delays due to crew constraints. We save anywhere from 1 to 5 million dollars on each major disruption. And with several of these disruptions in 2001 including the tropical storm Allison here in Houston and the tragic event of September 11th, the benefits go far beyond dollars. The improvement of our operations from passenger and crew perspectives is the most important benefit. In addition, it fits seamlessly into our operations and our overall strategy for dealing with disruptions. As we look into the future and an ever-changing environment, we always remember how important it is to remain competitive. Thanks to these technological advances which take full advantage of operations research, the future for crews and passengers is brighter than ever before."

## 1.2 General Description of Disruption Management

The case study in the previous section introduces the concept and philosophy of disruption management in the airline industry, demonstrates its remarkable value, and gives a brief account of the methodology and technology that are applied to put it into practice. We now formally lay out the framework, provide definitions to various terms, and present a general description for disruption management.

Disruptions caused by internal and external factors often cause a system to significantly deviate from its original plan, and severely affect its performance. Changes to the plan may originate from the various sources categorized below.

Changes in system environment

The environment in which the system is operating has changed unexpectedly. Such a change will affect the system's performance. For example, snowstorms may affect air and ground transportations; typhoons will severely impact logistics especially near harbor areas.

Unpredictable events

Spontaneous events unanticipated in the planning stage may severely impact the system. Examples include terrorist attacks, union strikes, power outages, etc.

Changes in system parameters

The parameters characterizing the system may have an unexpected change. For example, the market price in a supply chain may change for raw materials; the delivery time is changed from vendors.

Changes in availability of resources

The resources used in the system may become unavailable due to failure, quality reasons, and sicknesses. Examples include machine failures, resign of key personnel, etc.

New restrictions

New restrictions added to the system may make the original plan inferior or even infeasible. Examples include new government laws, new union contracts, new industry regulations, etc.

Uncertainties in system performance

Very often due to limited understanding of the system, its realized performance may fall below of our expectations. For example, the time for completing a task in a new project may be far from the estimated and thus may cause the delay of the entire project.

New considerations

New considerations that were not present in the planning phase must be handled properly and in a timely manner, or the system will bear cost and penalties. For example, in case of new customer orders or changed customer priorities if the system does not respond promptly and properly, it will lose its share of market.

In general, we denote the plan change as a "disruption" regardless of its cause and nature. Though some of the above issues can be described by certain probabilities and considered in an operational plan, due to the reasons we are to discuss later, the execution of the plan is still subject to real-time revision. In addition, disruptions also include situations that are very difficult or even impossible to anticipate such as a severe typhoon, the September 11th terrorist attack, and the massive blackouts in North America in August of 2003.

Disruption management can be formally stated as follows: At the beginning of a business cycle, an optimal or near-optimal operational plan is obtained by using certain optimization models and solution schemes. When such an operational plan is executed, disruptions may occur from time to time caused by internal and external uncertain factors. As a result, the original operational plan may not remain optimal, or even feasible. Consequently, we need to dynamically revise the original plan and obtain a new one that reflects the constraints and objectives of the evolved environment while minimizing the negative impact of the disruption. This process is referred to as disruption management.

## 1.3 Approaches to Uncertainties

In the past several decades, people have made great efforts to cope with uncertainties via different approaches. From the highest level, these approaches can be classified into two different stages: in-advance planning and real-time re-planning. The purpose of in-advance planning is to generate an optimal operational plan based on the estimate to future uncertainties. In real-time re-planning, the task is to revise the original plan in its execution period whenever needed.

*Contingency Planning*

Contingency planning is completely scenario-based. It uses a pre-allocated set of resources and a well-documented recipe to cope with disruptions. For example, to deal with SARS, the government may designate a select group of hospitals for accommodating SARS patients. Each time a suspect emerges, a pre-defined course of actions is taken for isolation, decontamination, examination, and treatment of the person and possibly everyone in contact. Thus, typical contingency planning consists of the following steps:

Step 1. During the planning stage, identify the critical scenarios. Here the criticality can be defined as probability of occurrence, the severity of potential impact, or combination of both.
Step 2. For each scenario, identify the options and associated effectiveness and costs.
Step 3. Based on available resources, choose the best option associated with the potential scenario.
Step 4. Acquire needed resources decided in the last step and place them in reserves. Clearly document the formula for dealing with each scenario. The formula will be closely followed upon occurrence of disruptions.

We immediately realize that such contingency planning can only handle a very limited set of well-understood disruptions, possibly localized, and evolving regularly based on certain rules. In a large system, the number of the possible future scenarios is vast, so it is very difficult to make a contingency plan for each scenario. When an unprepared scenario occurs, the performance of the system may worsen.

### *Stochastic Models*

A typical method of generating an operational plan within an uncertain environment is to use models based on stochastic process. In this way, a contingency plan or policy can be constructed that is optimal in terms of the average outcome. Ideally, a perfect contingency plan can be made that allows for all future possibilities. In other words, it has been perfectly planned for every future possibility. No real-time re-planning is needed. To do this with any measure of success, however, the precise probability distribution of future uncertainty must be known in advance, which is virtually impossible except in very simple cases. Moreover, even if the probability distribution is known, the form it takes may be too complicated to deal with, and thus some compromise would be needed, which could undermine the analysis. In any case, it is unavoidable that any operational plan generated based on the estimate to future uncertainty will have to be revised in its execution period when the uncertainties are resolved.

A typical operational plan or policy based on stochastic models contains the following steps.

Step 1. Build stochastic models to describe future uncertainty.
Step 2. Analyze the stochastic model and find the optimal policy so that the future output is optimized in terms of the average output.
Step 3. Execute the plan by taking the obtained policy for each scenario that occurs.

*Robust Optimization*

Robust optimization is another approach to handling uncertainty in the planning stage. In robust optimization, future uncertainties are modeled by a set of scenarios. The philosophy of robust optimization is to generate an operational plan that is "good" for most scenarios and acceptable for the worst scenario. The characteristic of robust optimization is that accurate probability distributions are not required; however, all possible scenarios must be specified. A decision can then be made so that it is optimal with respect to some criterion related to the worst-case outcome.

The purpose of robust optimization with regard to uncertainty is to keep the original plan intact no matter what happens. If a solution is chosen accordingly, the original plan will not cause an extremely inferior result. However, the solution may be too conservative if the worst-case scenario is associated with a very small probability. Moreover, it may still be difficult to state all possible scenarios when generating the original plan, even though the probability is not required. Therefore, in practice a robust solution may still be subject to change in its execution stage.

To illustrate robust planning in airline scheduling, suppose we know that a city is highly prone to snowstorms in the winter time. Then first of all, we should not choose the city as an airline hub during the network design phase unless it justifies other important criteria. We may also schedule flights with longer connection times passing the airport and allocate more reserve crews in the city during winter. Thus the worst scenario is considered. When a snowstorm hits, it will not affect a large scope of the network. Disruptions are absorbed locally.

A typical robust planning process is as follows:

Step 1. Identify the potential disruptive scenarios.
Step 2. Choose a robustness criterion appropriate for the decision maker. For example, one criterion defined in Kouvelis and Yu (1995) specifies the robustness as minimizing the maximum deviation from optimality under all possible scenarios.
Step 3. Incorporate the above information and measure in planning to generate a robust plan.
Step 4. Carry out the plan without change no matter what may happen in the future.

The advantage of a robust plan is that it can guarantee the performance of the system even when the worst scenario happens. However, the trade off is that a robust plan will sacrifice average performance to gain robustness against disruptions, especially when the probability of some disruptive events may be very small. Moreover, it is virtually impossible to enumerate all possible future disruptions in advance.

In general, an optimized operational plan can be generated based on either stochastic models or robust optimization models. From the viewpoint of practice, however, no perfect plan exists. Moreover, the execution of an operational plan usually involves some related plans made by other departments within an organization, but their interrelationship is usually neglected when the plans are generated. For example, the system for generating production planning may not take into account the difficulty in compiling a corresponding staffing assignment. The systems for generating production schedules and inventory policies tend to disregard the complexity of its consequent distribution issues. The system for manpower planning does not consider training scheduling and training resource use. The systems for planning decisions do not offer robust solutions needed for speedy recovery when there are disruptions. In all, this lack of integration among different related operational plans leads to inferior responsiveness to change and imposes costs throughout the enterprise. A disruption to one plan will also impact other related plans. Therefore, any operational plan would be subject to dynamic revision in its execution stage.

### *Disruption Management*

The concept of disruption management refers to the real time dynamic revision of an operational plan when disruptions occur. This is especially important in situations where an operational plan has to be published in advance, and its execution is subject to severe random disruptions. When a published operational plan is revised, there will be some deviation cost associated with the transition from the original plan to the new plan. The deviation cost can be a real dollar cost caused by raw material waste, or using on-call or reserved personnel; it can also mean the loss of the customers' goodwill for waiting and delay. To reduce such deviation costs, it is essential to take them into account when generating the new plan. In this sense, providing guidance for revising a predetermined plan is at least as important as making the plan itself.

### *Pure Rescheduling*

People have also been working on the research of various dynamic scheduling problems, on-line scheduling problems, and rescheduling models for decades. However, many current rescheduling problems have not been able to take into account the deviation costs in analysis. Thus we use pure rescheduling models to refer to the above research. In situations where there is no need to make and publish an operational schedule in advance, or a published schedule can be revised with little penalty, the pure rescheduling models can provide a solution to dynamically respond to the changing environment. However, when revision of a predetermined schedule is costly, failing to consider the deviation cost cannot provide a practically satisfactory solution.

We briefly compare the pros and cons of the various approaches below.

The contingency planning requires the listing of a finite set of scenarios. It can thus only handle this limited set of disruptions. If the real world event is not included in the set, then it may create havoc to the system. As possible disruptions are so unpredictable and very often cannot even be identified ahead of time, the limitation for the use of this approach is obvious.

Stochastic models are the most prevailing means of handling uncertainties, and have been approved to be effective in many situations. However, a successful application of stochastic models requires the knowledge of probability distribution, which is not always an easy job.

The robust planning approach also requires knowledge of the scenarios and possibly their probabilities of occurrence. The probability information on disruption scenarios is very difficult to acquire and its accuracy is often problematic. Thus a worst-case robustness measure is often used. However, such a measure emphasizes too much worst cases of small probability and thus makes the plan conservative.

Pure rescheduling focuses on reactive response to the changing environment, but it often ignores the possible deviation cost that may be incurred when changing the original plan, and thus only generates suboptimal solutions.

The major advantage of disruption management is its ability to cope with all disruptions without knowing what is going to unravel, and provide feasible and optimal solutions in real time.

## 1.4 General Guidance of Disruption Management

From our research and practice on disruption management, we have summarized some general guidelines and concepts for disruption management.

*Real-time optimization.* Disruption management is a real-time practice and often requires a quick solution when a disruption occurs. The original planning problem usually is regarded as a one-time effort, so it is practically acceptable if generating an optimal operational plan takes a dozens of minutes or hours, or even longer. However, when a disruption occurs, it is critical to immediately provide a resolution to the responsible personnel. For a large scale system, this is not always an easy job. Therefore, real-time optimization techniques should be developed.

*Deviation costs.* In handling disruptions, we need to correctly identify and quantify the deviation costs and take them into account. Without considering deviation costs, the recovery solution may contain too many undesirable changes to the plan and may be difficult or infeasible to implement the changes due to organizational, human, and other issues. As we have pointed out, the deviation costs may or may not be measured by a real dollar value. One of the roles of introducing the deviation costs is to force the revised plan to stay close to the original plan. In a large organization, many sub-systems are running interdependently. The change of one sub-system will impact other sub-systems. Such an impact is not always described by mathematical models. Therefore, we should make a very careful effort of constructing a balance between staying close to the original operational plan and reducing the cost of doing so.

*Multicriteria decision making.* A disruption management problem usually is modeled as a bi- or multi-criteria decision making problem. On the one hand, we have an original operational plan that has been optimized based on some criteria. Such criteria should still be taken into account when we are about to revise the original plan because they usually represent the goal of the operational plan, such as maximizing the profits. On the other hand, deviation costs should also be considered.

*Returning to the original plan.* In some circumstances, the revised plan may need to converge to the original plan, especially when the original plan is a long term repetitive plan, such as either a flight schedule in transportation or a contracted fixed reordering cycle in logistics. Usually, these plans are highly optimized. So when deviating from the plans, the system runs in a costly way. One of the goals of disruption management is to return to the original plan in a timely manner to reduce the cost and impact of a disruption.

*Disruption management time window.* When a disruption occurs, the decision maker may designate a time point by which the system should restore to its normal operation. This is referred to as the disruption management time window. By setting the time window, the impact of a disruption can be contained within a limited time period. This is necessary for the situations where it is important for the new plan to return to the original plan. We should also realize that there is a trade-off between recovery time and recovery cost. The sooner we would like to recover, the higher cost it will potentially incur.

*Multiple solutions.* In dealing with a disruption, it may be desired to generate multiple different high-quality solutions for the decision maker to review. This is important for several reasons such as having to consider multiple criteria, not being able to include all information in the optimization model, only near optimal solutions being generated in real time, and the coordination with other operational plans in the system.

*Partial solutions.* A partial solution refers to a solution that does not satisfy all existing constraints. A partial solution is not allowed in conventional optimization models but is an important concept and practice in disruption management. Facing a real time decision, the "buy-time" policy requires to generate a solution that can be put into execution immediately. Such a solution may violate some constraints that are less important or related to future time periods, which are left to be resolved gradually as time goes on.

## 1.5 An Overview of the Book

In writing this book, we have tried to make each chapter self-contained so that readers interested in some of the topics do not need to review the entire book. Meanwhile, reductant information in different chapters has been condensed as much as possible. In general, after reviewing the first two chapters for an overall description, the readers can read each following individual chapter separately. The contents of each chapter is briefed as follows.

In Chapter 2, we introduce several general models and methodologies for disruption management. A goal programming model is developed to address the concerns of multiple criteria, and the means of handling these multiple criteria. A scenario-based model is used to describe the disruption management as a dynamic process over time. Some local search approaches to solve disruption management problems are discussed.

In Chapters 3 and 4, we focus on disruption management for applications in airline operations. When a disruption occurs, the airline first needs to assess the disruption, then make revision to their flight and aircraft schedule. This will then disrupt their crew assignment and pairings, which, in turn, has to be repaired. We will address the flight and aircraft rescheduling in Chapter 3, and the crew rescheduling in Chapter 4.

Chapter 5 is concerned with disruption management for machine scheduling models. We present a systematic classification scheme for the research of disruption management for machine scheduling problems. We also discuss in detail problems where the original schedule is given by the shortest-processing-time rule, for both single and multiple machines.

Chapter 6 extends the results to logistics scheduling models. Different from traditional machine scheduling that is mainly concerned with sequencing jobs within a firm, logistics scheduling has to consider more issues between different locations such as transportation time and cost. So the disruption management becomes more important and complicated.

In Chapter 7, we consider the disruption management for continuous-time production and inventory management, in particular the EPQ/EOQ models. Both a single stage EPQ model and a two-stage series model are discussed.

In Chapter 8, we concentrate on disruption management for a discrete-time production and inventory model. The most interesting result there shows that under certain conditions, a greedy local search method can solve the problem to optimality with the least number of steps.

Chapter 9 is dedicated to disruption management for supply chain management. Specifically, we start with a two-player supply chain coordination model subject to demand disruptions. Then the results are extended to cases with more complicated demand functions and a one-supplier-two-retailer model.

In Chapter 10, we study disruption management for project scheduling. A project contains multiple activities which are intercorrelated by limited resource and precedence relationship. Mathematical models and solution schemes are proposed and tested.

## 1.6 Literature Review

The real-time operations control in the airline industry is one of the successful areas where the models and methodologies of disruption management have been applied. A lot of research results and applications are reported, for example, the flight rescheduling area including Teodorovic and Guberinic (1984), Jarrah *et al.*(1993), Teodorovic and Stojkovic (1995), Argüello *et al.*(1997), Cao and Kanafani (1997a, 1997b), Luo and Yu (1997), Yan and Lin (1997), Yan and Tu (1997), Thengvall *et al.*(2000, 2001), Filar *et al.*(2001), Bard *et al.*(2001), and crew rescheduling area including Wei *et al.*(1997), Lettovsky *et al.*(2000), Yu *et al.*(2003).

Machine scheduling and rescheduling is another area that people have been studying in order to respond to an unexpected disruptions. We refer to Aytug *et al.* (2001) for an extensive survey of rescheduling and executing production schedules while facing uncertainties. For work specifically addressing the deviation impact and cost of changing the original schedule, see Bean et al. (1991), Wu *et al.*(1993), Abumaizar and Svestka (1997), Akturk and Gorgulu (1999), Hall and Potts (2004).

Some efforts in discussing disruption and delays in project scheduling can be found in Howick and Eden (2001), Eden *et al.*(2002), and Howick (2003). The application of disruption management in other areas can be found in Clausen *et al.*(2001), and Sheffi (2003).

## References

Akturk, M.S. and E. Gorgulu (1999), Match-up scheduling under a machine breakdown, *European Journal of Operational Research*, 112, 81-97.

Argüello, M.F., J.F. Bard, and G. Yu (1997), A GRASP for aircraft routing in response to groundings and delays, *Journal of Combinatorial Optimization*, 5, 211-228.

Aytug, H., M.A. Lawley, K. McKay, S. Mohan and R. Uzsoy (2002), Executing production schedules in the face of uncertainties: A review and some future directions, *European Journal of Operational Research*, in press.

Bard, J.F., G. Yu and M.F. Argüello (2001), Optimizing aircraft routings in response to groundings and delays, *IIE Transactions on Operations Engineering*, 33(10), 931-947.

Bean, J.C., J.R. Birge, J. Mittenthal and C.E. Noon (1991), Matchup scheduling with multiple resources, release dates and disruptions, *Operations Research*, 39, 470-483.

Bethune, G. (1999), *From Worst to First: Behind the Scenes of Continental's Remarkable Comeback*, John Wiley & Sons.

Cao, J. and A. Kanafani (1997a), Real-time decision support for integration of airline flight cancellations and delays, part I: mathematical formulations, *Transportation Planning and Technology*, 20, 183-199.

Cao, J. and A. Kanafani (1997b), Real-time decision support for integration of airline flight cancellations and delays, part II: algorithms and computational experiments, *Transportation Planning and Technology*, 20, 201-217.

Clausen, J., J. Hansen, J. Larson and A. Larsen (2001), Disruption management, *ORMS Today*, 28, 40-43.

Eden, C., T. Williams, F. Ackermann, and S. Howick (2002), On the nature of disruption and delay (D&D) in major projects, *Journal of the Operational Research Society*, 51, 291-300.

Filar, J.A., P. Manyem and K. White (2001), How airlines and airports recover from schedule perturbations: A survey, *Annals of Operations Research*, 108, 315-333.

Hall, N.G. and C. Potts (2004), Recheduling for new orders, *Operations Research*, in press.

Howick, S. (2003), Using system dynamics to analyse disruption and delay in complex projects for litigation: Can the modelling purposes be met? *Journal of the Operational Research Society*, 54, 222-229.

Howick, S. and C. Eden (2001), The impact of disruption and delay when compressing large projects: Going for incentives? *Journal of the Operational Research Society*, 52, 26-34.

Jarrah, A.I.Z., G. Yu, N. Krishnamurthy and A. Rakshit (1993), A decision support framework for airline flight cancellations and delays, *Transportation Science*, 27, 266-280.

Kouvelis, P. and G. Yu (1996), *Robust Discrete Optimization and Its Applications*, Kluwer Academic Publishers, Boston.

Lettovsky, L., E.L. Johnson and G.L. Nemhauser (2000), Airline crew recovery, *Transportation Science*, 34, 337-348.

Luo, S.J. and G. Yu (1997), On the airline schedule perturbation problem caused by the ground delay program, *Transportation Science*, 31, 298-311.

Sheffi, Y. (2003), Supply chain disruption management, Presentation at the CMI Supply Chains Under Stress conference, Adastral Park, Ipswitch, United Kingdom.

Teodorovic, D. and S. Guberinic (1984), Optimal dispatching strategy on an airline network after a schedule perturbation, *European Journal of Operational Research*, 15, 178-182.

Teodorovic, D. and G. Stojkovic (1995), Model to reduce airline schedule disturbances, *Journal of Transportation Engineering*, 121, 324-331.

Thengvall, B.G., J.F. Bard and G. Yu (2000), Balancing user preferences for aircraft schedule recovery during irregular operations, *IIE Transactions on Operations Engineering*, 32, 181-193.

Thengvall, B.G., G. Yu and J.F. Bard (2001), Multiple fleet aircraft schedule recovery following hub closures, *Transportation Research Part A: Policy and Practice*, 35, 289-308.

Wei, G., G. Yu and M. Song (1997), Optimization model and algorithm for crew management during airline irregular operations, *Journal of Combinatorial Optimization*, 1, 305-321.

Yan, S. and C. Lin (1997), Airline scheduling for the temporary closure of airports, *Transportation Science*, 31, 72-82.

Yan, S. and Y.-P. Tu (1997), Multifleet routing and multistop flight scheduling for schedule perturbation, *European Journal of Operational Research*, 103, 155-169.

Yu, G., M. Argüello, M. Song, S. McCowan and A. White (2003), A new era for crew recovery at Continental Airlines, *Interfaces*, 33(1), 5-22.

Wu, S.-D., R.H. Storeer and P.C. Chang (1993), On machine rescheduling heuristics with efficiency and stability as criteria, *Computers & Operations Research*, 20, 1-14.