1

Introduction

1.1 THE ORIGINS OF OPERATIONS RESEARCH

Since the advent of the industrial revolution, the world has seen a remarkable growth in the size and complexity of organizations. The artisans' small shops of an earlier era have evolved into the billion-dollar corporations of today. An integral part of this revolutionary change has been a tremendous increase in the division of labor and segmentation of management responsibilities in these organizations. The results have been spectacular. However, along with its blessings, this increasing specialization has created new problems, problems that are still occurring in many organizations. One problem is a tendency for the many components of an organization to grow into relatively autonomous empires with their own goals and value systems, thereby losing sight of how their activities and objectives mesh with those of the overall organization. What is best for one component frequently is detrimental to another, so the components may end up working at cross purposes. A related problem is that as the complexity and specialization in an organization increase, it becomes more and more difficult to allocate the available resources to the various activities in a way that is most effective for the organization as a whole. These kinds of problems and the need to find a better way to solve them provided the environment for the emergence of **operations research** (commonly referred to as **OR**).

The roots of OR can be traced back many decades, when early attempts were made to use a scientific approach in the management of organizations. However, the beginning of the activity called *operations research* has generally been attributed to the military services early in World War II. Because of the war effort, there was an urgent need to allocate scarce resources to the various military operations and to the activities within each operation in an effective manner. Therefore, the British and then the U.S. military management called upon a large number of scientists to apply a scientific approach to dealing with this and other strategic and tactical problems. In effect, they were asked to do *research on* (military) *operations.* These teams of scientists were the first OR teams. By developing effective methods of using the new tool of radar, these teams were instrumental in winning the Air Battle of Britain. Through their research on how to better manage convoy and antisubmarine operations, they also played a major role in winning the Battle of the North Atlantic. Similar efforts assisted the Island Campaign in the Pacific.

When the war ended, the success of OR in the war effort spurred interest in applying OR outside the military as well. As the industrial boom following the war was running its course, the problems caused by the increasing complexity and specialization in organizations were again coming to the forefront. It was becoming apparent to a growing number of people, including business consultants who had served on or with the OR teams during the war, that these were basically the same problems that had been faced by the military but in a different context. By the early 1950s, these individuals had introduced the use of OR to a variety of organizations in business, industry, and government. The rapid spread of OR soon followed.

At least two other factors that played a key role in the rapid growth of OR during this period can be identified. One was the substantial progress that was made early in improving the techniques of OR. After the war, many of the scientists who had participated on OR teams or who had heard about this work were motivated to pursue research relevant to the field; important advancements in the state of the art resulted. A prime example is the *simplex method* for solving linear programming problems, developed by George Dantzig in 1947. Many of the standard tools of OR, such as linear programming, dynamic programming, queueing theory, and inventory theory, were relatively well developed before the end of the 1950s.

A second factor that gave great impetus to the growth of the field was the onslaught of the *computer revolution.* A large amount of computation is usually required to deal most effectively with the complex problems typically considered by OR. Doing this by hand would often be out of the question. Therefore, the development of electronic digital computers, with their ability to perform arithmetic calculations thousands or even millions of times faster than a human being can, was a tremendous boon to OR. A further boost came in the 1980s with the development of increasingly powerful personal computers accompanied by good software packages for doing OR. This brought the use of OR within the easy reach of much larger numbers of people. Today, literally millions of individuals have ready access to OR software. Consequently, a whole range of computers from mainframes to laptops now are being routinely used to solve OR problems.

1.2 THE NATURE OF OPERATIONS RESEARCH

As its name implies, operations research involves "research on operations." Thus, operations research is applied to problems that concern how to conduct and coordinate the *operations* (i.e., the *activities*) within an organization. The nature of the organization is essentially immaterial, and, in fact, OR has been applied extensively in such diverse areas as manufacturing, transportation, construction, telecommunications, financial planning, health care, the military, and public services, to name just a few. Therefore, the breadth of application is unusually wide.

The *research* part of the name means that operations research uses an approach that resembles the way research is conducted in established scientific fields. To a considerable extent, the *scientific method* is used to investigate the problem of concern. (In fact, the term *management science* sometimes is used as a synonym for operations research.) In particular, the process begins by carefully observing and formulating the problem, including gathering all relevant data. The next step is to construct a scientific (typically mathematical) model that attempts to abstract the essence of the real problem. It is then hypothesized that this model is a sufficiently precise representation of the essential features of the situation that the conclusions (solutions) obtained from the model are also valid for the real problem. Next, suitable experiments are conducted to test this hypothesis, modify it as needed, and eventually verify some form of the hypothesis. (This step is frequently referred to as *model validation.*) Thus, in a certain sense, operations research involves creative scientific research into the fundamental properties of operations. However, there is more to it than this. Specifically, OR is also concerned with the practical management of the organization. Therefore, to be successful, OR must also provide positive, understandable conclusions to the decision maker(s) when they are needed.

Still another characteristic of OR is its broad viewpoint. As implied in the preceding section, OR adopts an organizational point of view. Thus, it attempts to resolve the conflicts of interest among the components of the organization in a way that is best for the organization as a whole. This does not imply that the study of each problem must give explicit consideration to all aspects of the organization; rather, the objectives being sought must be consistent with those of the overall organization.

An additional characteristic is that OR frequently attempts to find a *best* solution (referred to as an *optimal* solution) for the problem under consideration. (We say *a* best instead of *the* best solution because there may be multiple solutions tied as best.) Rather than simply improving the status quo, the goal is to identify a best possible course of action. Although it must be interpreted carefully in terms of the practical needs of management, this "search for optimality" is an important theme in OR.

All these characteristics lead quite naturally to still another one. It is evident that no single individual should be expected to be an expert on all the many aspects of OR work or the problems typically considered; this would require a group of individuals having diverse backgrounds and skills. Therefore, when a full-fledged OR study of a new problem is undertaken, it is usually necessary to use a *team approach.* Such an OR team typically needs to include individuals who collectively are highly trained in mathematics, statistics and probability theory, economics, business administration, computer science, engineering and the physical sciences, the behavioral sciences, and the special techniques of OR. The team also needs to have the necessary experience and variety of skills to give appropriate consideration to the many ramifications of the problem throughout the organization.

1.3 THE IMPACT OF OPERATIONS RESEARCH

Operations research has had an impressive impact on improving the efficiency of numerous organizations around the world. In the process, OR has made a significant contribution to increasing the productivity of the economies of various countries. There now are a few dozen member countries in the International Federation of Operational Research Societies (IFORS), with each country having a national OR society. Both Europe and Asia have federations of OR societies to coordinate holding international conferences and publishing international journals in those continents.

It appears that the impact of OR will continue to grow. For example, according to the U.S. Bureau of Labor Statistics, OR currently is one of the fastest-growing career areas for U.S. college graduates.

To give you a better notion of the wide applicability of OR, we list some actual award-winning applications in Table 1.1. Note the diversity of organizations and applications in the first two columns. The curious reader can find a complete article describing each application in the January–February issue of *Interfaces* for the year cited in the third col-

TABLE 1.1 Some applications of operations research

Organization	Nature of Application	Year of Publication*	Related Chapters†	Annual Savings
The Netherlands Rijkswaterstaat	Develop national water management policy, including mix of new facilities, operating procedures, and pricing.	1985	2–8, 13, 22	$15 million
Monsanto Corp.	Optimize production operations in chemical plants to meet production targets with minimum cost.	1985	2, 12	$2 million
United Airlines	Schedule shift work at reservation offices and airports to meet customer needs with minimum cost.	1986	2–9, 12, 17, 18, 20	$6 million
Citgo Petroleum Corp.	Optimize refinery operations and the supply, distribution, and marketing of products.	1987	2–9, 20	$70 million
San Francisco Police Department	Optimally schedule and deploy police patrol officers with a computerized system.	1989	2–4, 12, 20	$11 million
Texaco, Inc.	Optimally blend available ingredients into gasoline products to meet quality and sales requirements.	1989	2, 13	$30 million
IBM	Integrate a national network of spare parts inventories to improve service support.	1990	2, 19, 22	$20 million +$250 million less inventory
Yellow Freight System, Inc.	Optimize the design of a national trucking network and the routing of shipments.	1992	2, 9, 13, 20, 22	$17.3 million
New Haven Health Department	Design an effective needle exchange program to combat the spread of HIV/AIDS.	1993	2	33% less HIV/AIDS
AT&T	Develop a PC-based system to guide business customers in designing their call centers.	1993	17, 18, 22	$750 million
Delta Airlines	Maximize the profit from assigning airplane types to over 2500 domestic flights.	1994	12	$100 million
Digital Equipment Corp.	Restructure the global supply chain of suppliers, plants, distribution centers, potential sites, and market areas.	1995	12	$800 million
China	Optimally select and schedule massive projects for meeting the country's future energy needs.	1995	12	$425 million
South African defense force	Optimally redesign the size and shape of the defense force and its weapons systems.	1997	12	$1.1 billion
Proctor and Gamble	Redesign the North American production and distribution system to reduce costs and improve speed to market.	1997	8	$200 million
Taco Bell	Optimally schedule employees to provide desired customer service at a minimum cost.	1998	12, 20, 22	$13 million
Hewlett-Packard	Redesign the sizes and locations of buffers in a printer production line to meet production goals.	1998	17, 18	$280 million more revenue

*Pertains to a January–February issue of *Interfaces* in which a complete article can be found describing the application.
†Refers to chapters in this book that describe the kinds of OR techniques used in the application.

umn of the table. The fourth column lists the chapters in *this* book that describe the kinds of OR techniques that were used in the application. (Note that many of the applications combine a variety of techniques.) The last column indicates that these applications typically resulted in annual savings in the millions (or even tens of millions) of dollars. Furthermore, additional benefits not recorded in the table (e.g., improved service to customers and better managerial control) sometimes were considered to be even more important than these financial benefits. (You will have an opportunity to investigate these less tangible benefits further in Probs. 1.3-1 and 1.3-2.)

Although most routine OR studies provide considerably more modest benefits than these award-winning applications, the figures in the rightmost column of Table 1.1 do accurately reflect the dramatic impact that large, well-designed OR studies occasionally can have.

We will briefly describe some of these applications in the next chapter, and then we present two in greater detail as case studies in Sec. 3.5.

1.4 ALGORITHMS AND OR COURSEWARE

An important part of this book is the presentation of the major **algorithms** (systematic solution procedures) of OR for solving certain types of problems. Some of these algorithms are amazingly efficient and are routinely used on problems involving hundreds or thousands of variables. You will be introduced to how these algorithms work and what makes them so efficient. You then will use these algorithms to solve a variety of problems on a computer. The CD-ROM called **OR Courseware** that accompanies the book will be a key tool for doing all this.

One special feature in your OR Courseware is a program called **OR Tutor.** This program is intended to be your personal tutor to help you learn the algorithms. It consists of many *demonstration examples* that display and explain the algorithms in action. These "demos" supplement the examples in the book.

In addition, your OR Courseware includes many *interactive routines* for executing the algorithms interactively in a convenient spreadsheet format. The computer does all the routine calculations while you focus on learning and executing the logic of the algorithm. You should find these interactive routines a very efficient and enlightening way of doing many of your homework problems.

In practice, the algorithms normally are executed by commercial software packages. We feel that it is important to acquaint students with the nature of these packages that they will be using after graduation. Therefore, your OR Courseware includes a wealth of material to introduce you to three particularly popular software packages described below. Together, these packages will enable you to solve nearly all the OR models encountered in this book very efficiently. We have added our own *automatic routines* to the OR Courseware only in a few cases where these packages are not applicable.

A very popular approach now is to use today's premier spreadsheet package, *Microsoft Excel,* to formulate small OR models in a spreadsheet format. The **Excel Solver** then is used to solve the models. Your OR Courseware includes a separate Excel file for nearly every chapter in this book. Each time a chapter presents an example that can be solved using Excel, the complete spreadsheet formulation and solution is given in that chapter's Excel file. For many of the models in the book, an *Excel template* also is provided that already includes all the equations necessary to solve the model. Some *Excel add-ins* also are included on the CD-ROM.

After many years, **LINDO** (and its companion modeling language **LINGO**) continues to be a dominant OR software package. Student versions of LINDO and LINGO now can be downloaded free from the Web. As for Excel, each time an example can be solved with this package, all the details are given in a LINGO/LINDO file for that chapter in your OR Courseware.

CPLEX is an elite state-of-the-art software package that is widely used for solving large and challenging OR problems. When dealing with such problems, it is common to also use a *modeling system* to efficiently formulate the mathematical model and enter it into the computer. **MPL** is a user-friendly modeling system that uses CPLEX as its main solver. A student version of MPL and CPLEX is available free by downloading it from the Web. For your convenience, we also have included this student version in your OR Courseware. Once again, all the examples that can be solved with this package are detailed in MPL/CPLEX files for the corresponding chapters in your OR Courseware.

We will further describe these three software packages and how to use them later (especially near the end of Chaps. 3 and 4). Appendix 1 also provides documentation for the OR Courseware, including OR Tutor.

To alert you to relevant material in OR Courseware, the end of each chapter from Chap. 3 onward has a list entitled *Learning Aids for This Chapter in Your OR Courseware.* As explained at the beginning of the problem section for each of these chapters, symbols also are placed to the left of each problem number or part where any of this material (including demonstration examples and interactive routines) can be helpful.

PROBLEMS

1.3-1. Select one of the applications of operations research listed in Table 1.1. Read the article describing the application in the January–February issue of *Interfaces* for the year indicated in the third column. Write a two-page summary of the application and the benefits (including nonfinancial benefits) it provided.

1.3-2. Select three of the applications of operations research listed in Table 1.1. Read the articles describing the applications in the January–February issue of *Interfaces* for the years indicated in the third column. For each one, write a one-page summary of the application and the benefits (including nonfinancial benefits) it provided.

2

Overview of the Operations Research Modeling Approach

The bulk of this book is devoted to the mathematical methods of operations research (OR). This is quite appropriate because these quantitative techniques form the main part of what is known about OR. However, it does not imply that practical OR studies are primarily mathematical exercises. As a matter of fact, the mathematical analysis often represents only a relatively small part of the total effort required. The purpose of this chapter is to place things into better perspective by describing all the major phases of a typical OR study.

One way of summarizing the usual (overlapping) phases of an OR study is the following:

1. Define the problem of interest and gather relevant data.
2. Formulate a mathematical model to represent the problem.
3. Develop a computer-based procedure for deriving solutions to the problem from the model.
4. Test the model and refine it as needed.
5. Prepare for the ongoing application of the model as prescribed by management.
6. Implement.

Each of these phases will be discussed in turn in the following sections.

Most of the award-winning OR studies introduced in Table 1.1 provide excellent examples of how to execute these phases well. We will intersperse snippets from these examples throughout the chapter, with references to invite your further reading.

2.1 DEFINING THE PROBLEM AND GATHERING DATA

In contrast to textbook examples, most practical problems encountered by OR teams are initially described to them in a vague, imprecise way. Therefore, the first order of business is to study the relevant system and develop a well-defined statement of the problem to be considered. This includes determining such things as the appropriate objectives, constraints on what can be done, interrelationships between the area to be studied and other areas of the organization, possible alternative courses of action, time limits for making a decision, and so on. This process of problem definition is a crucial one because it greatly affects how relevant the conclusions of the study will be. It is difficult to extract a “right” answer from the “wrong” problem!

The first thing to recognize is that an OR team is normally working in an *advisory capacity.* The team members are not just given a problem and told to solve it however they see fit. Instead, they are advising management (often one key decision maker). The team performs a detailed technical analysis of the problem and then presents recommendations to management. Frequently, the report to management will identify a number of alternatives that are particularly attractive under different assumptions or over a different range of values of some policy parameter that can be evaluated only by management (e.g., the tradeoff between *cost* and *benefits*). Management evaluates the study and its recommendations, takes into account a variety of intangible factors, and makes the final decision based on its best judgment. Consequently, it is vital for the OR team to get on the same wavelength as management, including identifying the "right" problem from management's viewpoint, and to build the support of management for the course that the study is taking.

Ascertaining the *appropriate objectives* is a very important aspect of problem definition. To do this, it is necessary first to identify the member (or members) of management who actually will be making the decisions concerning the system under study and then to probe into this individual's thinking regarding the pertinent objectives. (Involving the decision maker from the outset also is essential to build her or his support for the implementation of the study.)

By its nature, OR is concerned with the welfare of the *entire organization* rather than that of only certain of its components. An OR study seeks solutions that are optimal for the overall organization rather than suboptimal solutions that are best for only one component. Therefore, the objectives that are formulated ideally should be those of the entire organization. However, this is not always convenient. Many problems primarily concern only a portion of the organization, so the analysis would become unwieldy if the stated objectives were too general and if explicit consideration were given to all side effects on the rest of the organization. Instead, the objectives used in the study should be as specific as they can be while still encompassing the main goals of the decision maker and maintaining a reasonable degree of consistency with the higher-level objectives of the organization.

For profit-making organizations, one possible approach to circumventing the problem of suboptimization is to use *long-run profit maximization* (considering the time value of money) as the sole objective. The adjective *long-run* indicates that this objective provides the flexibility to consider activities that do not translate into profits *immediately* (e.g., research and development projects) but need to do so *eventually* in order to be worthwhile. This approach has considerable merit. This objective is specific enough to be used conveniently, and yet it seems to be broad enough to encompass the basic goal of profit-making organizations. In fact, some people believe that all other legitimate objectives can be translated into this one.

However, in actual practice, many profit-making organizations do not use this approach. A number of studies of U.S. corporations have found that management tends to adopt the goal of *satisfactory profits,* combined with *other objectives,* instead of focusing on long-run profit maximization. Typically, some of these *other* objectives might be to maintain stable profits, increase (or maintain) one's share of the market, provide for product diversification, maintain stable prices, improve worker morale, maintain family control of the business, and increase company prestige. Fulfilling these objectives might achieve long-run profit maximization, but the relationship may be sufficiently obscure that it may not be convenient to incorporate them all into this one objective.

Furthermore, there are additional considerations involving social responsibilities that are distinct from the profit motive. The five parties generally affected by a business firm located in a single country are (1) the *owners* (stockholders, etc.), who desire profits (dividends, stock appreciation, and so on); (2) the *employees,* who desire steady employment at reasonable wages; (3) the *customers,* who desire a reliable product at a reasonable price; (4) the *suppliers,* who desire integrity and a reasonable selling price for their goods; and (5) the *government* and hence the *nation,* which desire payment of fair taxes and consideration of the national interest. All five parties make essential contributions to the firm, and the firm should not be viewed as the exclusive servant of any one party for the exploitation of others. By the same token, international corporations acquire additional obligations to follow socially responsible practices. Therefore, while granting that management's prime responsibility is to make profits (which ultimately benefits all five parties), we note that its broader social responsibilities also must be recognized.

OR teams typically spend a surprisingly large amount of time *gathering relevant data* about the problem. Much data usually are needed both to gain an accurate understanding of the problem and to provide the needed input for the mathematical model being formulated in the next phase of study. Frequently, much of the needed data will not be available when the study begins, either because the information never has been kept or because what was kept is outdated or in the wrong form. Therefore, it often is necessary to install a new computer-based *management information system* to collect the necessary data on an ongoing basis and in the needed form. The OR team normally needs to enlist the assistance of various other key individuals in the organization to track down all the vital data. Even with this effort, much of the data may be quite "soft," i.e., rough estimates based only on educated guesses. Typically, an OR team will spend considerable time trying to improve the precision of the data and then will make do with the best that can be obtained.

Examples. An OR study done for the **San Francisco Police Department**[1] resulted in the development of a computerized system for optimally scheduling and deploying police patrol officers. The new system provided annual savings of \$11 million, an annual \$3 million increase in traffic citation revenues, and a 20 percent improvement in response times. In assessing the *appropriate objectives* for this study, three fundamental objectives were identified:

1. Maintain a high level of citizen safety.
2. Maintain a high level of officer morale.
3. Minimize the cost of operations.

To satisfy the first objective, the police department and city government jointly established a desired level of protection. The mathematical model then imposed the requirement that this level of protection be achieved. Similarly, the model imposed the requirement of balancing the workload equitably among officers in order to work toward the second objective. Finally, the third objective was incorporated by adopting the long-term goal of minimizing the number of officers needed to meet the first two objectives.

[1]P. E. Taylor and S. J. Huxley, "A Break from Tradition for the San Francisco Police: Patrol Officer Scheduling Using an Optimization-Based Decision Support System," *Interfaces,* **19**(1): 4–24, Jan.–Feb. 1989. See especially pp. 4–11.

The **Health Department of New Haven, Connecticut** used an OR team[1] to design an effective needle exchange program to combat the spread of the virus that causes AIDS (HIV), and succeeded in reducing the HIV infection rate among program clients by 33 percent. The key part of this study was an innovative *data collection program* to obtain the needed input for mathematical models of HIV transmission. This program involved complete tracking of *each* needle (and syringe), including the identity, location, and date for each person receiving the needle and each person returning the needle during an exchange, as well as testing whether the returned needle was HIV-positive or HIV-negative.

An OR study done for the **Citgo Petroleum Corporation**[2] optimized both refinery operations and the supply, distribution, and marketing of its products, thereby achieving a profit improvement of approximately $70 million per year. *Data collection* also played a key role in this study. The OR team held data requirement meetings with top Citgo management to ensure the eventual and continual quality of data. A state-of-the-art management database system was developed and installed on a mainframe computer. In cases where needed data did not exist, LOTUS 1-2-3 screens were created to help operations personnel input the data, and then the data from the personal computers (PCs) were uploaded to the mainframe computer. Before data was inputted to the mathematical model, a preloader program was used to check for data errors and inconsistencies. Initially, the preloader generated a paper log of error messages 1 inch thick! Eventually, the number of error and warning messages (indicating bad or questionable numbers) was reduced to less than 10 for each new run.

We will describe the overall Citgo study in much more detail in Sec. 3.5.

2.2 FORMULATING A MATHEMATICAL MODEL

After the decision maker's problem is defined, the next phase is to reformulate this problem in a form that is convenient for analysis. The conventional OR approach for doing this is to construct a mathematical model that represents the essence of the problem. Before discussing how to formulate such a model, we first explore the nature of models in general and of mathematical models in particular.

Models, or idealized representations, are an integral part of everyday life. Common examples include model airplanes, portraits, globes, and so on. Similarly, models play an important role in science and business, as illustrated by models of the atom, models of genetic structure, mathematical equations describing physical laws of motion or chemical reactions, graphs, organizational charts, and industrial accounting systems. Such models are invaluable for abstracting the essence of the subject of inquiry, showing interrelationships, and facilitating analysis.

[1]E. H. Kaplan and E. O'Keefe, "Let the Needles Do the Talking! Evaluating the New Haven Needle Exchange," *Interfaces,* **23**(1): 7–26, Jan.–Feb. 1993. See especially pp. 12–14.

[2]D. Klingman, N. Phillips, D. Steiger, R. Wirth, and W. Young, "The Challenges and Success Factors in Implementing an Integrated Products Planning System for Citgo," *Interfaces,* **16**(3): 1–19, May–June 1986. See especially pp. 11–14. Also see D. Klingman, N. Phillips, D. Steiger, and W. Young, "The Successful Deployment of Management Science throughout Citgo Petroleum Corporation," *Interfaces,* **17**(1): 4–25, Jan.–Feb. 1987. See especially pp. 13–15. This application will be described further in Sec. 3.5.

Mathematical models are also idealized representations, but they are expressed in terms of mathematical symbols and expressions. Such laws of physics as $F = ma$ and $E = mc^2$ are familiar examples. Similarly, the mathematical model of a business problem is the system of equations and related mathematical expressions that describe the essence of the problem. Thus, if there are n related quantifiable decisions to be made, they are represented as **decision variables** (say, $x_1, x_2, \ldots, x_n$) whose respective values are to be determined. The appropriate measure of performance (e.g., profit) is then expressed as a mathematical function of these decision variables (for example, $P = 3x_1 + 2x_2 + \cdots + 5x_n$). This function is called the **objective function.** Any restrictions on the values that can be assigned to these decision variables are also expressed mathematically, typically by means of inequalities or equations (for example, $x_1 + 3x_1x_2 + 2x_2 \leq 10$). Such mathematical expressions for the restrictions often are called **constraints.** The constants (namely, the coefficients and right-hand sides) in the constraints and the objective function are called the **parameters** of the model. The mathematical model might then say that the problem is to choose the values of the decision variables so as to maximize the objective function, subject to the specified constraints. Such a model, and minor variations of it, typifies the models used in OR.

Determining the appropriate values to assign to the parameters of the model (one value per parameter) is both a critical and a challenging part of the model-building process. In contrast to textbook problems where the numbers are given to you, determining parameter values for real problems requires *gathering relevant data.* As discussed in the preceding section, gathering accurate data frequently is difficult. Therefore, the value assigned to a parameter often is, of necessity, only a rough estimate. Because of the uncertainty about the true value of the parameter, it is important to analyze how the solution derived from the model would change (if at all) if the value assigned to the parameter were changed to other plausible values. This process is referred to as **sensitivity analysis,** as discussed further in the next section (and much of Chap. 6).

Although we refer to "the" mathematical model of a business problem, real problems normally don't have just a single "right" model. Section 2.4 will describe how the process of testing a model typically leads to a succession of models that provide better and better representations of the problem. It is even possible that two or more completely different types of models may be developed to help analyze the same problem.

You will see numerous examples of mathematical models throughout the remainder of this book. One particularly important type that is studied in the next several chapters is the **linear programming model,** where the mathematical functions appearing in both the objective function and the constraints are all linear functions. In the next chapter, specific linear programming models are constructed to fit such diverse problems as determining (1) the mix of products that maximizes profit, (2) the design of radiation therapy that effectively attacks a tumor while minimizing the damage to nearby healthy tissue, (3) the allocation of acreage to crops that maximizes total net return, and (4) the combination of pollution abatement methods that achieves air quality standards at minimum cost.

Mathematical models have many advantages over a verbal description of the problem. One advantage is that a mathematical model describes a problem much more concisely. This tends to make the overall structure of the problem more comprehensible, and it helps to reveal important cause-and-effect relationships. In this way, it indicates more clearly what additional data are relevant to the analysis. It also facilitates dealing with the problem in its entirety and considering all its interrelationships simultaneously. Finally, a mathematical model forms a bridge to the use of high-powered mathematical techniques and computers to analyze the problem. Indeed, packaged software for both personal computers and mainframe computers has become widely available for solving many mathematical models.

However, there are pitfalls to be avoided when you use mathematical models. Such a model is necessarily an abstract idealization of the problem, so approximations and simplifying assumptions generally are required if the model is to be *tractable* (capable of being solved). Therefore, care must be taken to ensure that the model remains a valid representation of the problem. The proper criterion for judging the validity of a model is whether the model predicts the relative effects of the alternative courses of action with sufficient accuracy to permit a sound decision. Consequently, it is not necessary to include unimportant details or factors that have approximately the same effect for all the alternative courses of action considered. It is not even necessary that the absolute magnitude of the measure of performance be approximately correct for the various alternatives, provided that their relative values (i.e., the differences between their values) are sufficiently precise. Thus, all that is required is that there be a high *correlation* between the prediction by the model and what would actually happen in the real world. To ascertain whether this requirement is satisfied, it is important to do considerable *testing* and consequent modifying of the model, which will be the subject of Sec. 2.4. Although this testing phase is placed later in the chapter, much of this *model validation* work actually is conducted during the model-building phase of the study to help guide the construction of the mathematical model.

In developing the model, a good approach is to begin with a very simple version and then move in evolutionary fashion toward more elaborate models that more nearly reflect the complexity of the real problem. This process of *model enrichment* continues only as long as the model remains tractable. The basic trade-off under constant consideration is between the *precision* and the *tractability* of the model. (See Selected Reference 6 for a detailed description of this process.)

A crucial step in formulating an OR model is the construction of the objective function. This requires developing a quantitative measure of performance relative to each of the decision maker's ultimate objectives that were identified while the problem was being defined. If there are multiple objectives, their respective measures commonly are then transformed and combined into a composite measure, called the **overall measure of performance.** This overall measure might be something tangible (e.g., profit) corresponding to a higher goal of the organization, or it might be abstract (e.g., utility). In the latter case, the task of developing this measure tends to be a complex one requiring a careful comparison of the objectives and their relative importance. After the overall measure of performance is developed, the objective function is then obtained by expressing this measure as a mathematical function of the decision variables. Alternatively, there also are methods for explicitly considering multiple objectives simultaneously, and one of these (goal programming) is discussed in Chap. 7.

Examples. An OR study done for **Monsanto Corp.**[1] was concerned with optimizing production operations in Monsanto's chemical plants to minimize the cost of meeting the target for the amount of a certain chemical product (maleic anhydride) to be produced in a given

[1] R. F. Boykin, "Optimizing Chemical Production at Monsanto," *Interfaces,* **15**(1): 88–95, Jan.–Feb. 1985. See especially pp. 92–93.

month. The decisions to be made are the dial setting for each of the catalytic reactors used to produce this product, where the setting determines both the amount produced and the cost of operating the reactor. The form of the resulting mathematical model is as follows:

Choose the values of the *decision variables* R_{ij}
$(i = 1, 2, \ldots, r;\ j = 1, 2, \ldots, s)$
so as to

$$\text{Minimize} \quad \sum_{i=1}^{r} \sum_{j=1}^{s} c_{ij} R_{ij},$$

subject to

$$\sum_{i=1}^{r} \sum_{j=1}^{s} p_{ij} R_{ij} \geq T$$

$$\sum_{j=1}^{s} R_{ij} = 1, \qquad \text{for } i = 1, 2, \ldots, r$$

$$R_{ij} = 0 \text{ or } 1,$$

where $R_{ij} = \begin{cases} 1 & \text{if reactor } i \text{ is operated at setting } j \\ 0 & \text{otherwise} \end{cases}$

c_{ij} = cost for reactor i at setting j
p_{ij} = production of reactor i at setting j
T = production target
r = number of reactors
s = number of settings (including off position)

The *objective function* for this model is $\Sigma\ \Sigma\ c_{ij}R_{ij}$. The *constraints* are given in the three lines below the objective function. The *parameters* are c_{ij}, p_{ij}, and T. For Monsanto's application, this model has over 1,000 *decision variables* R_{ij} (that is, $rs > 1{,}000$). Its use led to annual savings of approximately $2 million.

The Netherlands government agency responsible for water control and public works, the **Rijkswaterstaat,** commissioned a major OR study[1] to guide the development of a new national water management policy. The new policy saved hundreds of millions of dollars in investment expenditures and reduced agricultural damage by about $15 million per year, while decreasing thermal and algae pollution. Rather than formulating *one* mathematical model, this OR study developed a comprehensive, integrated system of 50 models! Furthermore, for some of the models, both simple and complex versions were developed. The simple version was used to gain basic insights, including trade-off analyses. The complex version then was used in the final rounds of the analysis or whenever greater accuracy or more detailed outputs were desired. The overall OR study directly involved over 125 person-years of effort (more than one-third in data gathering), created several dozen computer programs, and structured an enormous amount of data.

[1]B. F. Goeller and the PAWN team: "Planning the Netherlands' Water Resources," *Interfaces,* **15**(1): 3–33, Jan.–Feb. 1985. See especially pp. 7–18.

2.3 DERIVING SOLUTIONS FROM THE MODEL

After a mathematical model is formulated for the problem under consideration, the next phase in an OR study is to develop a procedure (usually a computer-based procedure) for deriving solutions to the problem from this model. You might think that this must be the major part of the study, but actually it is not in most cases. Sometimes, in fact, it is a relatively simple step, in which one of the standard **algorithms** (systematic solution procedures) of OR is applied on a computer by using one of a number of readily available software packages. For experienced OR practitioners, finding a solution is the fun part, whereas the real work comes in the preceding and following steps, including the *postoptimality analysis* discussed later in this section.

Since much of this book is devoted to the subject of how to obtain solutions for various important types of mathematical models, little needs to be said about it here. However, we do need to discuss the nature of such solutions.

A common theme in OR is the search for an **optimal,** or best, **solution.** Indeed, many procedures have been developed, and are presented in this book, for finding such solutions for certain kinds of problems. However, it needs to be recognized that these solutions are optimal only with respect to the model being used. Since the model necessarily is an idealized rather than an exact representation of the real problem, there cannot be any utopian guarantee that the optimal solution for the model will prove to be the best possible solution that could have been implemented for the real problem. There just are too many imponderables and uncertainties associated with real problems. However, if the model is well formulated and tested, the resulting solution should tend to be a good approximation to an ideal course of action for the real problem. Therefore, rather than be deluded into demanding the impossible, you should make the test of the practical success of an OR study hinge on whether it provides a better guide for action than can be obtained by other means.

Eminent management scientist and Nobel Laureate in economics Herbert Simon points out that **satisficing** is much more prevalent than optimizing in actual practice. In coining the term *satisficing* as a combination of the words *satisfactory* and *optimizing,* Simon is describing the tendency of managers to seek a solution that is "good enough" for the problem at hand. Rather than trying to develop an overall measure of performance to optimally reconcile conflicts between various desirable objectives (including well-established criteria for judging the performance of different segments of the organization), a more pragmatic approach may be used. Goals may be set to establish minimum satisfactory levels of performance in various areas, based perhaps on past levels of performance or on what the competition is achieving. If a solution is found that enables all these goals to be met, it is likely to be adopted without further ado. Such is the nature of satisficing.

The distinction between optimizing and satisficing reflects the difference between theory and the realities frequently faced in trying to implement that theory in practice. In the words of one of England's OR leaders, Samuel Eilon, "Optimizing is the science of the ultimate; satisficing is the art of the feasible."[1]

OR teams attempt to bring as much of the "science of the ultimate" as possible to the decision-making process. However, the successful team does so in full recognition of the

[1]S. Eilon, "Goals and Constraints in Decision-making," *Operational Research Quarterly,* **23:** 3–15, 1972—address given at the 1971 annual conference of the Canadian Operational Research Society.

overriding need of the decision maker to obtain a satisfactory guide for action in a reasonable period of time. Therefore, the goal of an OR study should be to conduct the study in an optimal manner, regardless of whether this involves finding an optimal solution for the model. Thus, in addition to pursuing the science of the ultimate, the team should also consider the cost of the study and the disadvantages of delaying its completion, and then attempt to maximize the net benefits resulting from the study. In recognition of this concept, OR teams occasionally use only **heuristic procedures** (i.e., intuitively designed procedures that do not guarantee an optimal solution) to find a good **suboptimal solution.** This is most often the case when the time or cost required to find an optimal solution for an adequate model of the problem would be very large. In recent years, great progress has been made in developing efficient and effective heuristic procedures (including so-called metaheuristics), so their use is continuing to grow.

The discussion thus far has implied that an OR study seeks to find only one solution, which may or may not be required to be optimal. In fact, this usually is not the case. An optimal solution for the original model may be far from ideal for the real problem, so additional analysis is needed. Therefore, **postoptimality analysis** (analysis done after finding an optimal solution) is a very important part of most OR studies. This analysis also is sometimes referred to as **what-if analysis** because it involves addressing some questions about *what* would happen to the optimal solution *if* different assumptions are made about future conditions. These questions often are raised by the managers who will be making the ultimate decisions rather than by the OR team.

The advent of powerful spreadsheet software now has frequently given spreadsheets a central role in conducting postoptimality analysis. One of the great strengths of a spreadsheet is the ease with which it can be used interactively by anyone, including managers, to see what happens to the optimal solution when changes are made to the model. This process of experimenting with changes in the model also can be very helpful in providing understanding of the behavior of the model and increasing confidence in its validity.

In part, postoptimality analysis involves conducting **sensitivity analysis** to determine which parameters of the model are most critical (the "sensitive parameters") in determining the solution. A common definition of *sensitive parameter* (used throughout this book) is the following.

> For a mathematical model with specified values for all its parameters, the model's **sensitive parameters** are the parameters whose value cannot be changed without changing the optimal solution.

Identifying the sensitive parameters is important, because this identifies the parameters whose value must be assigned with special care to avoid distorting the output of the model.

The value assigned to a parameter commonly is just an *estimate* of some quantity (e.g., unit profit) whose exact value will become known only after the solution has been implemented. Therefore, after the sensitive parameters are identified, special attention is given to estimating each one more closely, or at least its range of likely values. One then seeks a solution that remains a particularly good one for all the various combinations of likely values of the sensitive parameters.

If the solution is implemented on an ongoing basis, any later change in the value of a sensitive parameter immediately signals a need to change the solution.

In some cases, certain parameters of the model represent policy decisions (e.g., resource allocations). If so, there frequently is some flexibility in the values assigned to these parameters. Perhaps some can be increased by decreasing others. Postoptimality analysis includes the investigation of such trade-offs.

In conjunction with the study phase discussed in the next section (testing the model), postoptimality analysis also involves obtaining a sequence of solutions that comprises a series of improving approximations to the ideal course of action. Thus, the apparent weaknesses in the initial solution are used to suggest improvements in the model, its input data, and perhaps the solution procedure. A new solution is then obtained, and the cycle is repeated. This process continues until the improvements in the succeeding solutions become too small to warrant continuation. Even then, a number of alternative solutions (perhaps solutions that are optimal for one of several plausible versions of the model and its input data) may be presented to management for the final selection. As suggested in Sec. 2.1, this presentation of alternative solutions would normally be done whenever the final choice among these alternatives should be based on considerations that are best left to the judgment of management.

Example. Consider again the **Rijkswaterstaat** OR study of national water management policy for the Netherlands, introduced at the end of the preceding section. This study did not conclude by recommending just a single solution. Instead, a number of attractive alternatives were identified, analyzed, and compared. The final choice was left to the Dutch political process, culminating with approval by Parliament. *Sensitivity analysis* played a major role in this study. For example, certain parameters of the models represented environmental standards. Sensitivity analysis included assessing the impact on water management problems if the values of these parameters were changed from the current environmental standards to other reasonable values. Sensitivity analysis also was used to assess the impact of changing the assumptions of the models, e.g., the assumption on the effect of future international treaties on the amount of pollution entering the Netherlands. A variety of *scenarios* (e.g., an extremely dry year and an extremely wet year) also were analyzed, with appropriate probabilities assigned.

2.4 TESTING THE MODEL

Developing a large mathematical model is analogous in some ways to developing a large computer program. When the first version of the computer program is completed, it inevitably contains many bugs. The program must be thoroughly tested to try to find and correct as many bugs as possible. Eventually, after a long succession of improved programs, the programmer (or programming team) concludes that the current program now is generally giving reasonably valid results. Although some minor bugs undoubtedly remain hidden in the program (and may never be detected), the major bugs have been sufficiently eliminated that the program now can be reliably used.

Similarly, the first version of a large mathematical model inevitably contains many flaws. Some relevant factors or interrelationships undoubtedly have not been incorporated into the model, and some parameters undoubtedly have not been estimated correctly. This is inevitable, given the difficulty of communicating and understanding all the aspects and subtleties of a complex operational problem as well as the difficulty of collecting reliable data. Therefore, before you use the model, it must be thoroughly tested to try to identify and correct as many flaws as possible. Eventually, after a long succession of improved models, the OR team concludes that the current model now is giving reasonably valid results. Although some minor flaws undoubtedly remain hidden in the model (and may never be detected), the major flaws have been sufficiently eliminated that the model now can be reliably used.

This process of testing and improving a model to increase its validity is commonly referred to as **model validation.**

It is difficult to describe how model validation is done, because the process depends greatly on the nature of the problem being considered and the model being used. However, we make a few general comments, and then we give some examples. (See Selected Reference 2 for a detailed discussion.)

Since the OR team may spend months developing all the detailed pieces of the model, it is easy to "lose the forest for the trees." Therefore, after the details ("the trees") of the initial version of the model are completed, a good way to begin model validation is to take a fresh look at the overall model ("the forest") to check for obvious errors or oversights. The group doing this review preferably should include at least one individual who did not participate in the formulation of the model. Reexamining the definition of the problem and comparing it with the model may help to reveal mistakes. It is also useful to make sure that all the mathematical expressions are *dimensionally consistent* in the units used. Additional insight into the validity of the model can sometimes be obtained by varying the values of the parameters and/or the decision variables and checking to see whether the output from the model behaves in a plausible manner. This is often especially revealing when the parameters or variables are assigned extreme values near their maxima or minima.

A more systematic approach to testing the model is to use a **retrospective test.** When it is applicable, this test involves using historical data to reconstruct the past and then determining how well the model and the resulting solution would have performed if they had been used. Comparing the effectiveness of this hypothetical performance with what actually happened then indicates whether using this model tends to yield a significant improvement over current practice. It may also indicate areas where the model has shortcomings and requires modifications. Furthermore, by using alternative solutions from the model and estimating their hypothetical historical performances, considerable evidence can be gathered regarding how well the model predicts the relative effects of alternative courses of actions.

On the other hand, a disadvantage of retrospective testing is that it uses the same data that guided the formulation of the model. The crucial question is whether the past is truly representative of the future. If it is not, then the model might perform quite differently in the future than it would have in the past.

To circumvent this disadvantage of retrospective testing, it is sometimes useful to continue the status quo temporarily. This provides new data that were not available when the model was constructed. These data are then used in the same ways as those described here to evaluate the model.

Documenting the process used for model validation is important. This helps to increase confidence in the model for subsequent users. Furthermore, if concerns arise in the future about the model, this documentation will be helpful in diagnosing where problems may lie.

Examples. Consider once again the **Rijkswaterstaat** OR study of national water management policy for the Netherlands, discussed at the end of Secs. 2.2 and 2.3. The process of model validation in this case had three main parts. First, the OR team checked the general behavior of the models by checking whether the results from each model moved in reasonable ways when changes were made in the values of the model parameters. Second, retrospective testing was done. Third, a careful technical review of the models, methodology, and results was conducted by individuals unaffiliated with the project, including Dutch experts. This process led to a number of important new insights and improvements in the models.

Many new insights also were gleaned during the model validation phase of the OR study for the **Citgo Petroleum Corp.,** discussed at the end of Sec. 2.1. In this case, the model of refinery operations was tested by collecting the actual inputs and outputs of the refinery for a series of months, using these inputs to fix the model inputs, and then comparing the model outputs with the actual refinery outputs. The process of properly calibrating and recalibrating the model was a lengthy one, but ultimately led to routine use of the model to provide critical decision information. As already mentioned in Sec. 2.1, the validation and correction of input data for the models also played an important role in this study.

Our next example concerns an OR study done for **IBM**[1] to integrate its national network of spare-parts inventories to improve service support for IBM's customers. This study resulted in a new inventory system that improved customer service while reducing the value of IBM's inventories by over $250 million and saving an additional $20 million per year through improved operational efficiency. A particularly interesting aspect of the model validation phase of this study was the way that *future users* of the inventory system were incorporated into the testing process. Because these future users (IBM managers in functional areas responsible for implementation of the inventory system) were skeptical about the system being developed, representatives were appointed to a *user team* to serve as advisers to the OR team. After a preliminary version of the new system had been developed (based on a multiechelon inventory model), a *preimplementation test* of the system was conducted. Extensive feedback from the user team led to major improvements in the proposed system.

2.5 PREPARING TO APPLY THE MODEL

What happens after the testing phase has been completed and an acceptable model has been developed? If the model is to be used repeatedly, the next step is to install a well-documented *system* for applying the model as prescribed by management. This system will include the model, solution procedure (including postoptimality analysis), and oper-

[1]M. Cohen, P. V. Kamesam, P. Kleindorfer, H. Lee, and A. Tekerian, "Optimizer: IBM's Multi-Echelon Inventory System for Managing Service Logistics," *Interfaces,* **20**(1): 65–82, Jan.–Feb. 1990. See especially pp. 73–76. This application will be described further in Sec. 19.8.

ating procedures for implementation. Then, even as personnel changes, the system can be called on at regular intervals to provide a specific numerical solution.

This system usually is *computer-based.* In fact, a considerable number of computer programs often need to be used and integrated. *Databases* and *management information systems* may provide up-to-date input for the model each time it is used, in which case interface programs are needed. After a solution procedure (another program) is applied to the model, additional computer programs may trigger the implementation of the results automatically. In other cases, an *interactive* computer-based system called a **decision support system** is installed to help managers use data and models to support (rather than replace) their decision making as needed. Another program may generate *managerial reports* (in the language of management) that interpret the output of the model and its implications for application.

In major OR studies, several months (or longer) may be required to develop, test, and install this computer system. Part of this effort involves developing and implementing a process for maintaining the system throughout its future use. As conditions change over time, this process should modify the computer system (including the model) accordingly.

Examples. The **IBM** OR study introduced at the end of Sec. 2.4 provides a good example of a particularly large computer system for applying a model. The system developed, called *Optimizer,* provides optimal control of service levels and spare-parts inventories throughout IBM's U.S. parts distribution network, which includes two central automated warehouses, dozens of field distribution centers and parts stations, and many thousands of outside locations. The parts inventory maintained in this network is valued in the billions of dollars. Optimizer consists of four major modules. A forecasting system module contains a few programs for estimating the failure rates of individual types of parts. A data delivery system module consists of approximately 100 programs that process over 15 gigabytes of data to provide the input for the model. A decision system module then solves the model on a weekly basis to optimize control of the inventories. The fourth module includes six programs that integrate Optimizer into IBM's Parts Inventory Management System (PIMS). PIMS is a sophisticated information and control system that contains millions of lines of code.

Our next example also involves a large computer system for applying a model to control operations over a national network. This system, called *SYSNET,* was developed as the result of an OR study done for **Yellow Freight System, Inc.**[1] Yellow Freight annually handles over 15 million shipments by motor carrier over a network of 630 terminals throughout the United States. SYSNET is used to optimize both the routing of shipments and the design of the network. Because SYSNET requires extensive information about freight flows and forecasts, transportation and handling costs, and so on, a major part of the OR study involved integrating SYSNET into the corporate management information system. This integration enabled periodic updating of all the input for the model. The implementation of SYSNET resulted in annual savings of approximately $17.3 million as well as improved service to customers.

[1]J. W. Braklow, W. W. Graham, S. M. Hassler, K. E. Peck, and W. B. Powell, "Interactive Optimization Improves Service and Performance for Yellow Freight System," *Interfaces,* **22**(1): 147–172, Jan.–Feb. 1992. See especially p. 163.

Our next example illustrates a *decision support system.* A system of this type was developed for **Texaco**[1] to help plan and schedule its blending operations at its various refineries. Called *OMEGA* (Optimization Method for the Estimation of Gasoline Attributes), it is an *interactive* system based on a nonlinear optimization model that is implemented on both personal computers and larger computers. Input data can be entered either manually or by interfacing with refinery databases. The user has considerable flexibility in choosing an objective function and constraints to fit the current situation as well as in asking a series of *what-if questions* (i.e., questions about *what* would happen *if* the assumed conditions change). OMEGA is maintained centrally by Texaco's information technology department, which enables constant updating to reflect new government regulations, other business changes, and changes in refinery operations. The implementation of OMEGA is credited with annual savings of more than $30 million as well as improved planning, quality control, and marketing information.

2.6 IMPLEMENTATION

After a system is developed for applying the model, the last phase of an OR study is to implement this system as prescribed by management. This phase is a critical one because it is here, and only here, that the benefits of the study are reaped. Therefore, it is important for the OR team to participate in launching this phase, both to make sure that model solutions are accurately translated to an operating procedure and to rectify any flaws in the solutions that are then uncovered.

The success of the implementation phase depends a great deal upon the support of both top management and operating management. The OR team is much more likely to gain this support if it has kept management well informed and encouraged management's active guidance throughout the course of the study. Good communications help to ensure that the study accomplishes what management wanted and so deserves implementation. They also give management a greater sense of ownership of the study, which encourages their support for implementation.

The implementation phase involves several steps. First, the OR team gives operating management a careful explanation of the new system to be adopted and how it relates to operating realities. Next, these two parties share the responsibility for developing the procedures required to put this system into operation. Operating management then sees that a detailed indoctrination is given to the personnel involved, and the new course of action is initiated. If successful, the new system may be used for years to come. With this in mind, the OR team monitors the initial experience with the course of action taken and seeks to identify any modifications that should be made in the future.

Throughout the entire period during which the new system is being used, it is important to continue to obtain feedback on how well the system is working and whether the assumptions of the model continue to be satisfied. When significant deviations from the original assumptions occur, the model should be revisited to determine if any modifications should be made in the system. The postoptimality analysis done earlier (as described in Sec. 2.3) can be helpful in guiding this review process.

[1]C. W. DeWitt, L. S. Lasdon, A. D. Waren, D. A. Brenner, and S. A. Melhem, "OMEGA: An Improved Gasoline Blending System for Texaco," *Interfaces,* **19**(1): 85–101, Jan.–Feb. 1989. See especially pp. 93–95.

Upon culmination of a study, it is appropriate for the OR team to *document* its methodology clearly and accurately enough so that the work is *reproducible. Replicability* should be part of the professional ethical code of the operations researcher. This condition is especially crucial when controversial public policy issues are being studied.

Examples. This last point about *documenting* an OR study is illustrated by the Rijkswaterstaat study of national water management policy for the Netherlands discussed at the end of Secs. 2.2, 2.3, and 2.4. Management wanted unusually thorough and extensive documentation, both to support the new policy and to use in training new analysts or in performing new studies. Requiring several years to complete, this documentation aggregated 4000 single-spaced pages and 21 volumes!

Our next example concerns the **IBM** OR study discussed at the end of Secs. 2.4 and 2.5. Careful planning was required to implement the complex Optimizer system for controlling IBM's national network of spare-parts inventories. Three factors proved to be especially important in achieving a successful implementation. As discussed in Sec. 2.4, the first was the inclusion of a *user team* (consisting of operational managers) as advisers to the OR team throughout the study. By the time of the implementation phase, these operational managers had a strong sense of ownership and so had become ardent supporters for installing Optimizer in their functional areas. A second success factor was a very extensive *user acceptance test* whereby users could identify problem areas that needed rectifying prior to full implementation. The third key was that the new system was *phased in gradually,* with careful testing at each phase, so the major bugs could be eliminated before the system went live nationally.

Our final example concerns **Yellow Freight's** SYSNET system for routing shipments over a national network, as described at the end of the preceding section. In this case, there were four key elements to the implementation process. The first was selling the concept to upper management. This was successfully done through validating the accuracy of the cost model and then holding *interactive sessions* for upper management that demonstrated the effectiveness of the system. The second element was the development of an implementation strategy for gradually phasing in the new system while identifying and eliminating its flaws. The third involved working closely with operational managers to install the system properly, provide the needed support tools, train the personnel who will use the system, and convince them of the usefulness of the system. The final key element was the provision of management incentives and enforcement for the effective implementation of the system.

2.7 CONCLUSIONS

Although the remainder of this book focuses primarily on *constructing* and *solving* mathematical models, in this chapter we have tried to emphasize that this constitutes only a portion of the overall process involved in conducting a typical OR study. The other phases described here also are very important to the success of the study. Try to keep in perspective the role of the model and the solution procedure in the overall process as you move through the subsequent chapters. Then, after gaining a deeper understanding of mathematical models, we suggest that you plan to return to review this chapter again in order to further sharpen this perspective.

OR is closely intertwined with the use of computers. In the early years, these generally were mainframe computers, but now personal computers and workstations are being widely used to solve OR models.

In concluding this discussion of the major phases of an OR study, it should be emphasized that there are many exceptions to the "rules" prescribed in this chapter. By its very nature, OR requires considerable ingenuity and innovation, so it is impossible to write down any standard procedure that should always be followed by OR teams. Rather, the preceding description may be viewed as a model that roughly represents how successful OR studies are conducted.

SELECTED REFERENCES

1. Fortuin, L., P. van Beek, and L. van Wassenhove (eds.): *OR at wORk: Practical Experiences of Operational Research,* Taylor & Francis, Bristol, PA, 1996.
2. Gass, S. I.: "Decision-Aiding Models: Validation, Assessment, and Related Issues for Policy Analysis," *Operations Research,* **31:** 603–631, 1983.
3. Gass, S. I.: "Model World: Danger, Beware the User as Modeler," *Interfaces,* **20**(3): 60–64, May–June 1990.
4. Hall, R. W.: "What's So Scientific about MS/OR?" *Interfaces,* **15**(2): 40–45, March–April 1985.
5. Miser, H. J.: "The Easy Chair: Observation and Experimentation," *Interfaces,* **19**(5): 23–30, Sept.–Oct. 1989.
6. Morris, W. T.: "On the Art of Modeling," *Management Science,* **13:** B707–717, 1967.
7. Murthy, D. N. P., N. W. Page, and E. Y. Rodin: *Mathematical Modeling: A Tool for Problem Solving in Engineering, Physical, Biological and Social Sciences,* Pergamon Press, Oxford, England, 1990.
8. Simon, H. A.: "Prediction and Prescription in Systems Modeling," *Operations Research,* **38:** 7–14, 1990.
9. Tilanus, C. B., O. B. DeGans, and J. K. Lenstra (eds.): *Quantitative Methods in Management: Case Studies of Failures and Successes,* Wiley, New York, 1986.
10. Williams, H. P.: *Model Building in Mathematical Programming,* 3d ed., Wiley, New York, 1990.

PROBLEMS

2.1-1. Read the article footnoted in Sec. 2.1 that describes an OR study done for the San Francisco Police Department.
(a) Summarize the background that led to undertaking this study.
(b) Define part of the problem being addressed by identifying the six directives for the scheduling system to be developed.
(c) Describe how the needed data were gathered.
(d) List the various tangible and intangible benefits that resulted from the study.

2.1-2. Read the article footnoted in Sec. 2.1 that describes an OR study done for the Health Department of New Haven, Connecticut.
(a) Summarize the background that led to undertaking this study.
(b) Outline the system developed to track and test each needle and syringe in order to gather the needed data.
(c) Summarize the initial results from this tracking and testing system.
(d) Describe the impact and potential impact of this study on public policy.

2.2-1. Read the article footnoted in Sec. 2.2 that describes an OR study done for the Rijkswaterstaat of the Netherlands. (Focus especially on pp. 3–20 and 30–32.)
(a) Summarize the background that led to undertaking this study.
(b) Summarize the purpose of each of the five mathematical models described on pp. 10–18.