# OPERATIONS RESEARCH CENTER

working paper



MASSACHUSETTS INSTITUTE OF TECHNOLOGY

# Computer-Aided Dispatch System As A Decision Making Tool in Public and Private Sectors

by

I-Jen Lee
Richard C. Larson

OR 129-84 December 1984

## ABSTRACT

We describe in detail seven distinct areas in both public and private sectors in which a real-time computer-aided dispatch system is applicable to the allocation of scarce resources. Characteristics of a real-time computer-aided dispatch system are identified. Monetary savings and/or non-monetary advantages are presented to justify the implementation of a computer-aided dispatch system.

TABLE OF CONTENTS

# I INTRODUCTION

This paper surveys a variety of areas where the so-called real-time computer-aided dispatch (CAD) system has been applied or can be applied to the allocation of scarce resources. A real-time Computer-aided dispatch system is a man/machine interface between dispatchers and on-line digital computers. The use of the system is to facilitate decision-making on the real-time deployment of resources.

Unlike an off-line system in which information on the current state of the system is ignored and relevant decisions are made separately, a real-time system takes into account knowledge of the current state of the system and can help a dispatcher make a series of decisions simultaneously or instantaneously. Therefore, a real-time planning algorithm (i.e., computer implementable procedures) should be able to tell us quickly what to do when a certain event has occurred.

## II MOTIVATION FOR BUILDING A REAL-TIME CAD SYSTEM

Although applications of real-time CAD systems cover a wide range of areas in both the public and private sectors, the driving force for all the developments in real-time CAD systems essentially results from the growth of computer technology, which includes increased computational capabilities and declining cost of computers. As the availability and use of computers increase, people begin to notice the benefits that might be generated from the operation of computers in the management of scarce resources.

In general, the use of computers in a CAD system can be classified into two categories. The first category is the automation of the operational process, such as customer request processing, producing management information reports, and record keeping. The second category is the automation of the dispatching process, which provides algorithmic decision support to dispatchers. The distinction between the two is that automation of the operational process simply replaces previous manual operations, whereas automation of the dispatching process utilizes CAD data to help dispatchers allocate scarce resources more efficiently. Although the automation of the operational process supplies, perhaps, more accurate information and faster processing than manual operations, the major benefits to be realized in a CAD system lie in the automation of the dispatching process. Two examples of benefits due to automation of the dispatching process are increased productivity and faster response time.

Since allocating resources efficiently is the main concern in the automation of the dispatching process, techniques from operations research are used to develop algorithmic procedures. Therefore, the success of a real-time CAD system depends not only on the continual growth of computer technology, but also on the creative and successful application of operations research techniques.

## III AN OVERVIEW OF EXISTING OR POTENTIAL CAD SYSTEMS

In this section, we will consider some already existent as well as potential applications of real-time CAD systems. Specifically, in each application area, we will discuss, when applicable, the following aspects: (a) the problem of concern, (b) the existing CAD system and its impact, and (c) further research.

### 1. Urban Police Services

The primary functions of the "dispatching center" or "communication center" of police departments is receiving calls from the public and dispatching police patrol units. Many police researchers have considered the dispatching function as the "brain" of the police patrol force, since the police dispatchers' decisions on minute-to-minute utilization of all radio dispatchable patrol units determine the performance of the entire police response system, based on performance measures such as the utilization of each police patrol unit, the response time to every call for service, the amount of directed and/or preventive patrol in each neighborhood by time of delay, and citizen satisfaction with the police response system.

In the report, "A National Assessment of Police Command, Control, and Communications systems,"[3] the authors point out that there are now more than 80 police computer-aided dispatch systems throughout the country which are operational or in the process of being implemented. A typical CAD system, according to the same authors, assists the police telephone operator, sometimes called a complaint clerk, in verifying incident-related information such as addresses and checking to see if certain information such as incident address and incident type has been entered. The computer then advances the information to the appropriate dispatcher. The dispatcher, upon studying the location, priority and other aspects of the call for service, again through the aid of a computer, selects one or more appropriate available police patrol cars to respond to the call for service. When the service is completed, the police officer in the dispatched car informs the dispatcher by the two-way radio of availability for another assignment.

Note that the current CAD systems do not assist in decision-making of either complaint clerks or dispatchers. Instead, the CAD systems are programmed to quickly categorize and manipulate files, generate lists, and organize backlogs.

In addition to the fiscal difficulties many police departments now face, including increasing demand, a decreasing number of personnel and declining financial support, several relevant police findings that will be described below strongly indicate the need to fully utilize on-line computers in CAD systems. Some of the findings, summarized in an MIT project proposal [11], are the following:

1. Rapid police response is important only for approximately 10 to 15% of all calls for service.

2. Citizen satisfaction with police response time depends much more critically on expectation of response time rather than the reality of response time.

3. By careful training of police call-takers and prioritization of calls for service, it is possible for police departments to engage in so-called management of demand. In other words, instead of the traditional reply, "A police car will be there as soon as possible,". The complaint clerk may take certain reports on the phone or ask the caller to give the reports in person or schedule on-scene police response to occur later in the day or perhaps even during the next day.

Motivated by the above factors, a research project titled, "Proactive Real-Time Management of Scarce Resources," has been conducted at the MIT Operations Research Center under the supervision of Professor Richard C. Larson. The objective of the research is to apply operations research techniques to the development of algorithms in order for CAD systems to (1) assist in prioritization of various types of calls, (2) provide guidance to the dispatcher as to when to enter calls in queue versus immediately dispatching them, and (3) provide estimates of various delays that might occur.

## 2. Dial-a-Ride Services

Unlike the police dispatching problem where little research has been done, the dial-a-ride problem has received considerable attention. In the dial-a-ride problem, customers call a dispatcher requesting service. Each customer specifies a distinct pickup and delivery point and, sometimes, a desired time for pickup or delivery as well. If all customers request immediate service, then the routing and scheduling is done in real-time and the problem is called the dynamic or real-time dial-a-ride problem. If all customers call in advance, then the demand for services is known before routing and scheduling is done and the problem is called the static or subscriber dial-a-ride problem.

There are many applications of dial-a-ride problems and their extensions. Some examples are campus escort services after darkness, shared cab rides, package delivery, and, perhaps more importantly, the routing and scheduling of vehicles for the elderly and handicapped. Although the United States, Canada, and Europe have computerized some of these and other so-called paratransit operations[14], Bodin et al. point out in their paper[2] that only a few of these communities actually route and schedule their vehicles by computer and most of the computerized operations are concerned with reservation and management information systems.

For our purposes, only the real-time dial-a-ride problem is discussed. In such a demand responsive system, the route and schedule of each vehicle at the time of a new customer's request are known. At this time, some of the earlier customers have been delivered to their destinations and, hence, are no longer considered in the problem. The remaining earlier customers are either on board their assigned vehicles waiting to be delivered to their destination or are still waiting to be picked up. The problem is to determine the assignment of the new customer to a vehicle and the new route and schedule for the vehicle this new customer is assigned to.

As one might expect, an intuitive way to solve this problem is to assign the new customer to the vehicle that is closest to the origin or pick-up point of this new customer. In fact, this is the rule both the Senior House in Dorchester and the MIT Campus Police have adopted. Most of the work on this problem has been carried out by Psaraftis, Odoni, Wilson and their colleagues at MIT. According to Psaraftis[15], demand-responsive transportation systems have been in operation in several metropolitan areas of the United States as well as abroad. Examples are Rochester, New York; Ann Arbor, Michigan; San Jose, California; Tokyo, Japan. Due to lack of information on the evaluations of these computerized transportation systems at Ann Arbor, San Jose, and Tokyo, in the remaining section, we shall only describe briefly the algorithm that is implemented in Rochester, and its evaluations.

The basic idea of the assignment algorithm in Rochester Dial-A-Ride Demonstration Project is to select the best way to incorporate the new customer into the existing set of scheduled vehicle activities by use of an insertion method as soon as the customer requests service. All possible combinations of insertions of pickup and delivery stops for this customer into the routes of all available vehicles are investigated. The combination which minimizes a certain objective function is chosen. The objective function consists of three terms: (1) disutility of service to the new customer, (2) marginal disutility to earlier customers, and (3) marginal system resources committed to serve the new customer. Instead of deriving a theoretical disutility function, Wilson uses a linear, approximated objective function. For details, see references [21] and [22]. Since there exists a great amount of uncertainty in vehicle and driver activities, Wilson further develops an automatic reassignment capability (i.e., the ability to systematically reconsider previous assignments) as part of the overall control procedure. Bodin et al. summarizes the features of such a procedure, which was later implemented in Rochester, as follows:

(i) Every 20 minutes, the system is examined. All customers in the system who have not as yet been picked up by a vehicle are reassigned to another vehicle. The reassignment goes as follows. First, customers assigned to vehicle 1 are reassigned, then vehicle 2, etc. The order in which customers from a vehicle are reassigned corresponds to the times when they are assigned to that vehicle. All customers on a vehicle or already delivered are not considered for reassignment.

(ii) Whenever a driver makes a stop out of sequence or whenever a vehicle is more than 15 minutes late in carrying out its route, all customers waiting to be served by that vehicle are considered for reassignment.

The goal of computer dispatching, according to Wilson, is to decrease cost per customer carried and/or to improve the quality of service provided. However, as Wilson explains, there are some difficulties in using either criterion as a system performance measure. First, since computer control for the Rochester control system incurs higher control costs than manual dispatching, a decrease in cost per customer requires an increase in vehicle productivities, which in turn, requires either a reduction in the number of vehicles operating while serving the same demand or an increase in the number of customers riding the same number of vehicles. It is, in practice, difficult to decrease the number of vehicles due to the long-term commitment of drivers to service. So, a reduction in cost per customer depends on the ability to improve ridership. But improving ridership depends on the flexibility of demand with respect to service quality. Thus, service quality is chosen to be a system performance. However, a serious problem arises, i.e., no specific measures of service quality, as perceived by the customer, have yet been established. Therefore, even though a set of service quality measures are proposed by analysts, one cannot be sure that these measures are correct and hence, no overall system performance can be identified.

Second, there are many variations in an uncontrolled real world that make the before-and-after analysis of computed control difficult. Some examples of the variations are availability of service (e.g., vehicle breakdowns, etc.), changes in operational policies (e.g., area size served) and exogenous factors (e.g., full crises).

In an attempt to make some general conclusion about what the Rochester demonstration has shown about computerized dispatching of demand responsive transit systems, Wilson has defined three components of service quality: wait time, travel time, and pickup time deviation (i.e., difference between promised pickup time and actual pickup time). Moreover, given a large sample of customers for whom these measures of service quality were obtained, Wilson uses the following three statistics to characterize overall performance in terms of each component of service quality: mean of total sample, standard deviation of total sample and standard deviation of daily mean. A comparison of computer and manual dispatching in Irondequoit, Rochester's oldest suburb, was carried out using the above system performance measures. The details of the results are presented in chapter 4 of Wilson's report[22] and we do not wish to repeat them here. In general, the Rochester demonstration has shown that computerized dispatching can have a significant positive effect on some elements of service quality provided to customers. In particular, wait time and service reliability have been much improved as a result of the automation of the dispatching process. However, some other elements of service quality have been adversely affected. For example, manual dispatching appears to provide more consistent ride times for immediate customers.

In short, the Rochester demonstration does not reveal whether or not the automation of demand responsive operation increases vehicle productivity and, therefore, is cost effective. It, however, suggests further investigations of this question. Finally, it is important to note that, above all, the Rochester demonstration has produced a working computerized demand-responsive dispatch and control package. Thus, the package can readily be used elsewhere with only minor software conversion.

## 3. Vehicle Routing and Scheduling

We now consider, separately, the experiences of two large chemical companies, E.I. DuPont, Inc. and Air Products and Chemicals, Inc., in their applications of operations research techniques to their problems of dispatching their delivery fleet for consumer goods.

**A. The Local Delivery Problem.** Before we explain the problems DuPont faces, it is necessary to describe briefly its distribution system for its consumable products. The Clinical System Division of DuPont markets to hospitals and other medical institutions worldwide a machine that automates many of the routine tests made for patients in medical laboratories. This machine is called the Automatic Clinical Analyzer and is denoted by aca. Any consumable products of a chemical nature (called the "consumables") are required to operate an aca. These consumables, available only from DuPont, must be refrigerated and are shipped directly to the customers from the Region Distribution Centers.

There are three different ways to ship the consumables to the customers. They are called DRD mode, non-DRD mode and transshipment mode. The DRD mode which stands for Direct Refrigerated Delivery, involves delivery by one of the medium size, flat-bed, refrigerated trucks, called DRD trucks, that drive weekly loops of several dozen customers. The non-DRD mode uses an outside carrier such as air freight or motor freight. The transshipment mode is to transship goods from a DRD truck to an outside carrier at a point closer to the customer than the truck's loading site. One kind of decision that DuPont must make is the assignment of customers to delivery mode, DRD trucks, and delivery loops.

Another kind of decision is concerned with so-called satellite terminals. Truck terminals are located either at a Regional Distribution Center or a place remote from a Regional Distribution Center. A remote truck terminal is called satellite. Originally, DuPont had only one satellite terminal in

Houston. The successful operation of Houston satellite terminal raised the following questions for the managers. If new satellites are added, where should they be located? Should the Houston satellite be expanded and moved? If more satellites are created, should some Regional Distribution Centers be eliminated?

Since DuPont only delivers the consumables to the customers once a month, it has agreed to deliver them on the same day each month. In order to satisfy these rigid schedules, reassignment of customers to loops is usually avoided and new customers are added to the existing loops. However, as more customers are added to the current loops, the loops become more inefficient. Therefore, DuPont righteousness the DRD loops twice a year; meanwhile, a few customers are reassigned to a new delivery mode.

Due to the rapid increase in the number of coustomers (it was estimated to be 40% per year, which caused the delivery cost to exceed $1 million annually, DuPont became increasingly concerned over the efficiency of the redesign process which was performed manually. As a result of a study conducted by Fisher, Jaikumar and their colleagues, DuPont now employs a state-of-the-art computerized vehicle routing package called the "Real Time Optimizer for Vehicle Routing" (ROVER) to create loops that minimize the cost of full, driver time, and vehicle depreciation, subject to truck capacity constraints, timing constraints, etc.

ROVER uses a sophisticated embedded optimization heuristic to determine the assignment of customers to each vehicle and the order in which the customers should be delivered to minimize the total travel cost. This optimization heuristic first finds an assignment of customers to vehicles by solving a generalized assignment problem with a linear objective function that approximates delivery cost. Given this assignment, it then applies the traveling salesman problem heuristic described in Lin and Kernigham [12] to determine a delivery sequence for the customers assigned to each vehicle. The generalized assignment problem is solved by using the Lagrangian relaxation method, in which the multipliers are found by applying a multiplier adjustment method given in Fisher [6]. Details of the generalized assignment approach to vehicle routing can be found [8].

It is interesting to note that a color graphics interface for ROVER was also developed to aid the schedule, or dispatcher in visualizing the progress of the program and in adjusting initial conditions to produce a better solution. The purpose of graphical analysis, as Fisher, Greenfield and Jaikumar explained in [7], is to exploit the pattern recognition capabilities of the human mind and yet use the computer to provide the images and perform necessary computation. Specifically, one point at which a color graphic interface is useful is in determining the linear objective function that approximates delivery cost for the generalized assignment problem. The decision support system is a pioneering work in its marriage of advanced optimization methods with state-of-the-art computer interface technology.

Finally, to answer the questions concerning satellite terminals, a list of alternative satellite scenarios was generated by inspection of a plot of the current loops and suggestions from management. Each alternative was investigated by reassigning customers and using ROVER to produce revised loops. The proposed new loops were sent to drivers and then refined to remove the problems identified by the drivers. The final solution is said to reduce delivery costs by approximately 15%.

B. **Inventory Delivery Problem.** Air Products and Chemicals, Inc. is a major manufacturer in the industrial gas industry. Its primary products are liquid oxygen and liquid nitrogen. These products are stored in large tanks and distributed in cryogenic bulk tankers to industrial users and hospitals. Since the manufacturing costs are generally the same for different suppliers, competition in the industry depends primarily on distribution efficiencies. Air Products and Chemicals, Inc., therefore, became interested in the implementation of a computerized scheduling system.

The distribution problems of Air Products and Chemicals, Inc. are similar to those of DuPont's in that we are given a fleet of vehicles with certain known capacities that delivers a bulk product stored at a central depot. However, because of the bulk nature of the product, the customer demand is specified differently. Instead of a fixed known amount of demand, a lower limit and an upper limit on the amount of product to be delivered during the planning horizon are given. A value per unit of product is rewarded to the ending inventory. The problem is to determine the assignment of customers to vehicles, the delivery sequence of the customers assigned to each vehicle as well as the number of deliveries to make for each customer and the amount of product to provide on each delivery so that the value of product delivered to all customers less delivery costs incurred is maximized.

It turns out that this problem -- inventory management combined with vehicle scheduling and dispatching (Fisher and Jaikumar call it the "Inventory Delivery Problem") can be formulated as a very large mix-integer program (up to 8000,000 variables and 200,000 constraints) which is beyond the current state of the art in mix-integer programming. (Typical problems involve, according to Bell, et al. [1], several hundred customers and about 20 trucks.) Fortunately, because of the special structure of the program with a sparse coefficient matrix, the mix-integer program can be decomposed into a set of knapsack problems -- one for each vehicle and is solved implicitly, again, by ROVER. The solutions produced by ROVER were proven to be within one-half to two percent of optimality in a modest amount of computing time. A complete description of the algorithm procedure for solving the mix-integer program of this problem is presented in [10].

ROVER is part of the vehicle scheduling system that has been implemented on a computer at Air Products corporate headquarters. Schedulers have access to the system over the phone lines on CRT terminals at the plants. ROVER is run daily and a typical run involves 500,000 variables and 150,000 constraints requiring 1 to 5 minutes of AMDAHL V8 commuter time.

Before we turn to another application, two other features of the vehicle scheduling system deserve our attention. First, unlike DuPont's experience where the distance and travel time between any pair of customers was validated to be adequately approximated by using the common approach of straight line distances and then adjusting by some factor, Air Products and Chemicals, Inc. had to adopt a different approach: a computerized network representation of the road system in the United States, along with a shortest route algorithm described in Glover and Klingman [9] produced the desired input data for ROVER. Second, a user-friendly interactive interface was developed to allow the scheduler to query and update input data for ROVER. Second, a user-friendly interactive interface was developed to allow the scheduler to query and update input for ROVER and to change the schedule as desired via a schedule change module. This feature is particularly important, since it enables the system to response to contingencies quickly. Moreover, it allows the schedulers who have developed intuition, through their experiences, in creating opportunities for improving productivities to utilize their judgments.

Since its first implementation at Wharton, New Jersey in October 1981, the vehicle scheduling system is reported to save between six to ten percent of operating costs, which, according to Jaikumar, Fisher, Greenfield and Kedia [10], will amount to several million dollars of savings annually once the system has been implemented throughout the United States.

Finally, other domains of applications of the inventory delivery problem, as Jaikumar, et al. suggest, include the delivery of petroleum products to gas stations and the delivery of liquid propane. In addition, some production planning problems with sequence dependent setup costs can be modeled as an inventory delivery problem if we identify a route with a production schedule for a family of products, the amount of product delivered with the amount of a product produced, and a truck with a specific machine.

## 4. Train Dispatching

Another area where a real-time computer-aided dispatch system has been proven to be beneficial is in train dispatching. Dispatching trains is a complex and demanding process. Normally, in an eight hour shift, a train dispatcher will control about 20 to 30 trains, most of which operate over single tracks. Thus, the dispatcher must arrange opposing trains to meet at passing sidings. He also must safely coordinate movements of roadway maintenance gangs, signal maintenance crews, etc. In addition, federal law requires that the dispatcher maintain a "train sheet" which records train operating and delay statistics. The train dispatcher also coordinates with other dispatchers, exchanging information about the operation of his territory.

Due to the openings of a large freight yard facility at Sheffield and a coal loading facility near Sheffield, the northern portion of Southern Railway Alabama Division, (i.e., the routes extending from Atlanta through Birmingham) to Sheffield, became heavily congested. The Southern Railway Company (now part of Norfolk Southern Corporation) installed a centralized traffic control (CTC), which provides a failsafe system of signals and switches in the field, controlled centrally by the dispatcher on an electronic display board. In addition to the CTC installation, information systems for yard and terminal operations were already in use at many locations on the railroad. These factors, perhaps, provided the motivation for the operations research staff of Southern Railway Company to investigate further computer assistance for the dispatcher.

Computer assistance was expected to yield two improvements -- in the automation of data collection and in the planning operations. The massive amount of division level information that was previously recorded manually was computer coded and integrated with other information systems. An on-line predictive planning aid was developed to assist dispatcher in dealing with operational problems, such as arranging opposing trains to meet at passing sidings. A mini-computer based simulator was built which emulated a centralized-traffic-control office environment and allowed the users to explore real-life scenarios. A so-called meet/pass planning algorithm which assisted dispatchers in making optimal routing decision under certain objective was incorporated in the simulator. The meet/pass planning algorithm, which is described in Sauder and Westerman's appendix [19], considers all potential meets of opposing trains and the locations (sidings) where these meets might be resolved, and then choose the route that gives rise to minimum total train delay. The total train delay is the sum of a certain function of the total train delay for each train. The total train delay for a train i consists of two terms: the delay, $W_i$, that would permit train i to reach its destination within a predefined scheduled run time and the delay, $y_i$, that exceeds this predefined scheduled run time, multiplied by train i's priority. The authors incorporated trains' priorities into the second term in order to promote adherence to schedule, especially for higher priority trains. A branch and bound enumeration technique is used to generate all feasible meet solutions. A feasible solution requires that all trains reach their destinations within 20 hours and that the current solution's train delay be less than those of previously enumerated solutions.

Operations with and without computer-aided planning were simulated. It turned out that without computer-aided planning, the total train delay resulting from the simulated operations was slightly less than that from manual operation. With computer-aided planning, however, the total train delay showed an improvement from simulated operation by 34 percent. These positive results encouraged the conversion of the simulator to an on-line environment by interfacing the mini-computers with the CTC system. A mini-computer based information system with on-line optimal route planning capability was installed and finally used to assist dispatchers on the northern portion of Southern Railway's Alabama Division in 1980.

After two years of operation, it was reported that train delay has been reduced by more than 15 percent, which translates into annual savings of $316,000. Besides monetary savings, the superior routing strategy and track management of the dispatching system provides additional track time for various types of maintenance crews. Furthermore, using the meet/pass planning capability, the dispatcher can now determine the best route that maximizes on-track working time while minimizing the train delay. This dispatching system is being expanded to other Southern Railway operating divisions. The ultimate goal is to move from the division-level computer-aided dispatching system to a system-wide optimizing operation.

5. Maintenance Operations

Maintenance is a potential area where a real-time computer-aided dispatch system might improve operational efficiency and generate savings. Some examples of maintenance operations are dispatching road repair crews to fix potholes which may occur from a Poisson distribution or from other probabilistic distributions and dispatching public housing maintenance crews to fix problems such as broken windows, leaking faucets, etc. that might occur in the apartments. Maintenance operations have the same type of management problems as police patrol operations. To illustrate this, let us consider the maintenance operations in public housing.

The public housing authority, funded by the federal government and state governments, provides public housing for low income families. The authority's duties include not only assigning applicants to houses but also sending maintenance crews to fix problems in the rented apartments. Up to now, the scheduling of the maintenance operations have been done manually.

Typically, when a problem arises, the tenants call up the development office of the public housing authority to report the incident and request services. The offices record the type of incident and the time and place where the incident occurred and then dispatch an appropriate service if one is available. If no server is available, the problem usually sits around and queues up until it gets assigned to someone.

The similarities between the problem of dispatching service crews in public housing's maintenance operations and the problem of dispatching police patrol units are apparent. In both application areas, "customers" (i.e., citizen or applicants to public housing) request services through telephone calls, a call-taker records the incident-related information and sends it to a dispatcher, the dispatcher responds to requests by checking the priority of the call, the current status of "servers" (maintenance crews or police patrol cars) and other aspects of the calls for service and finally selects a server to respond or places the call in queue. In addition, both application areas share the same characteristic that the on-scene service times dominate the travel times. However, in public housing's maintenance operations, the service times are about three hours on average, whereas in urban police services, the service times range from 20 minutes to an hour. Despite the difference in the length of service times, it seems plausible that, as in urban police services, a computer-aided dispatch system with an on-line optimal planning capability will improve productivity in public housing's maintenance operations.

Finally, similar applications extend also to some emergency response stations, or services such as fire departments, ambulances, snow emergency services, and power line repair services.

6. Oil Spill Cleanups

Oil spill cleanups are perhaps a less known emergency service to the general populace. Oil spills, however, are of great concern to governments, as well as to the ocean transportation and oil exploration and production industries. Besides massive catastrophic spills which happen rarely, small and moderate spills occur on a daily basis from various sources such as operational discharges from tankers, vessel collisions, etc. An oil spill cleanup is an emergency action that must be taken to prevent further pollution of the sea and coastline once a spill occurs.

Although many studies regarding oil spills have been done, very few of them have addressed the complex decision-making process in the oil spill cleanup. In an attempt to understand this process and to assist decision-makers, or, in the case of the Coast Guard, the On Scene Coordinator, in dealing with the problems of oil spills, a computer aided model which integrates all parts of a spill response system and incorporates quantitative descriptors of system performance as well as decision-making techniques was developed by a research group at MIT. The model, called the Oil Spill Model, which was implemented in an interactive model on a VAX 780/AA Computer at MIT, consists of three components: inputs, submodels and outputs. A description of the entire model, which will not be included here, can be found in [16] or [17].

One of the submodels, called the tactical model, was developed to assist decision-makers in the optimal allocation of resources for cleaning up a specific spill after its occurrence was known. In contrast to the so-called operational model where questions regarding a more detailed oil spill response action such as the spatial allocation of cleanup resources are concerned, the tactical model addresses the aggregate level of oil recovery capability and answers questions such as what equipment should be dispatched to the spill site. Both tactical and operational decisions concern actions in response to a specific spill and therefore should be dynamically adjusted during a spill as changes occur in oil outflow rate and weather condition. In the remaining section, we shall only discuss the tactical decision problem and its application to the ARGO MERCHANT Oil spill which occurred in December 1976.

The tactical decision problem, defined in [18], can be summarized as follows. An oil spill of known oil type and quantity occurs at some specific time and some known geographical location. The spill is reported to the responsible authorities at a given time. The authorities must determine the response tactic that minimizes a weighted combination of spill-specific cleanup and damage costs, given the following assumptions: (1) a variety of oil spill cleanup equipment sets (i.e., an integrated, sufficiently equipped and self-contained package) are located at known sites and are available for dispatch, (2) the performance of each equipment set is a known function of the prevailing weather conditions, (3) future weather conditions and spill outflow are known with certainty, and (4) the costs of using each equipment set and the damage costs of the oil spill are known. Note that, in order to simplify the analysis, the authors assumed, in their first phase of the development of the model, that the problem is deterministic, even though the inputs to the problem including factors such as spill outflow rate or weather condition are inherently stochastic in the real world. Later, we shall explain a dynamic version of the tactical decision problem proposed by the authors.

To solve the tactical decision problem described above, the authors employed an approximate version of a dynamic programming algorithm ("Approximate" in the sense that a surrogate representation of the state variable, equipment deployment on scene, was used) within which a series of 0-1 knapsack problems were solved repeatedly. The authors quantified the worst-case performance of this approximate version and claimed that this procedure would produce solutions very close to optimality under realistic inputs. For a detailed discussion of the formulation of the problem as well as its solution techniques, see [18].

We now consider an application of the MIT Oil Spill Model. In particular, we shall only discuss the experience related to the tactical aspect of the application.

The MIT Oil Spill Mode was first tested with the ARGO MERCHANT oil spill, which occurred off the coast of Massachusetts with a release of 7.7 million gallons of heavy crude oil. The

U.S. Coast Guard responded with a massive mobilization of cleanup equipment, manpower and scientific support. However, no oil was offloaded from the tanker, ARGO MERCHANT, or removed from the sea. In fact, as a result of prevailing northwesterly winds, the oil drifted offshore. Thus, the only damage cost was the lost oil.

What interested the MIT research team, the authors claimed, was not the actual account of the oil spill, but an investigation of "what if" questions concerning that spill. Some examples of the kinds of questions of interest are: "What would have happened if winds in the ARGO MERCHANT oil spill were blowing the opposite direction?", or "What if the spill had occurred in the midst of the summer tourist season and winds carried it onto the New England coast?" These two and other related questions were analyzed from the viewpoints of oil spill trajectory, equipment dispatching, cleanup costs and damage assessment. A complete description of these runs will not be presented here, but can be found in [14]. Instead, we would like to point out a potential use of the MIT Oil Spill Model, which is suggested by the capability of answering "what if" questions. This potential use is the ability to assist the On Scene Coordinator in real-time optimal response to spills. We illustrate this point below.

Recall that, for simplicity, the authors assumed that the tactical decision problem is deterministic, even though the input variables to this problem are all stochastic in the real world. One way to take into account the stochastic nature of the input data is to run the tactical model in a static mode for a number of alternative scenarios whose probabilities have been determined a priori. One would then obtain a probability distribution of the consequences of those scenarios and then choose the response tactic that is most consistent with the decision-maker's risk preference structure. Another way to account for the stochastic nature of the input data is to run the tactical model by first entering the user's best guess (i.e., expected) forecasts for the entire duration of the spill event, and then dynamically update inputs and rerun the algorithm whenever a significant forecast revision is made.

Finally, we should note that both methods described above are not genuine stochastic extensions to the tactical decision problem (or the tactical model). Although stochastic modeling of both slick trajectories and weather conditions have been attempted by many researchers, for instance, Dietzel et al. [5] and Devanney [4], none of these methods have yet been incorporated, or successfully applied by the MIT research team or others.

## 7. Refuse Disposal

In this section, we consider another computer-aided dispatch system that was implemented in New York City. The problem of concern is refuse disposal. In New York City, refuse is dumped into large, floating containers that are located at either Marine Transfer Stations (MTS) or some temporary barge docks called Staging Areas. The barges, which have no power of self-propulsion, must be transported through waterways by marine towing vehicles, called tugs, from MTS to the ultimate disposal site, Fresh Kills Landfill (FKL). After barges are unloaded at FKL, tugs must transport these unloaded barges (called "light barges") back to marine transfer stations. The dispatcher is concerned with the pick-up and delivery assignments for tugs leaving FKL.

In 1982, a study on the barge fleet operations of New York City was carried out by the ENFORTH Corporation. As a result of the study, a Barge Operations System Simulator (BOSS) was developed as a planning model for the New York City Department of Sanitation. A complete description of this simulator, as well as a user's manual for the BOSS computer package can be found in [13]. In this paper, we shall only discuss the dispatching algorithm within BOSS.

The basic idea behind the dispatching algorithm is to select the assignment of tugs to MTSs that minimizes the tonnage of refuse deferred throughout the marine operations system. A simple example perhaps can best illustrate the barge assignment process. Assume that a tug is waiting to leave Fresh Kills Landfill with a certain number of light barges. The dispatcher first predicts the time for which each marine transfer station becomes blocked (i.e., no more refuse can be accommodated by the barges) and, hence, requires a tug towing light barges to visit. Such time, called the tug block time, always exists, since there are only a finite number of barges at each MTS. The dispatcher considers all possible sequences of dispatches to the marine transfer stations. For each sequence, the times of arrival of dispatched tugs are then estimated. Suppose that, for a particular sequence, each marine transfer station is visited before blockage can occur, except, say, MTS3 and MTS4. Then the total estimated deferred tonnage for this sequence is the sum of the mean amount of refuse delivered between the tug block time and the dispatched tug arrival time. If for the current run of the simulation, the sequence examined yields a total deferred tonnage that is strictly less than the previous minimum, then this sequence is considered the best thus far. At the end of this simulation series, the best solution is chosen and then implemented.

One should note that the above example is a simple illustration of the dispatching algorithm within the BOSS computer package. There are other features of the dispatching algorithm that are worth mentioning, but are not essential to our discussion under the scope of this paper. The interested readers, again, are referred to [13].

Before we leave this section, we would like to point out a difficulty associated with the simulation. Because of the complexity of barge fleet operations, the BOSS model, a simulation model as opposed to an analytical model, was developed to provide the means for observing the consequences of all manners of change in the marine system. Consequently, the BOSS model suffers from a common problem that occurs in all simulation models. The problem is the following: Each run of the simulation can be considered as a "computer experiment," representing what would happen if the real system had been operated under some chosen configuration for a specified period of time. However, the measured value (e.g., mean queueing delay at FKL) obtained from a simulation run is usually not the same as the long-term actual average value that would be found if one studied the system for a longer period of time. In spite of this difficulty, the authors report that, for most of the system configurations tested, the performance of the model was very good. Thus, the performance of the model has justified its use.

## IV CONCLUSION

We have identified seven distinct areas where a real-time computer-aided dispatch system has either been successfully applied or has shown its potential for productivity improvement. One should note, however, that this list of application areas is not exhaustive.

In general, any application of real-time CAD systems can be characterized as follows. First, there exists a fleet of vehicles to be dispatched. The fleet may have certain type, size and capacities. Second, there is a spatial distribution of the occurrences of customers by customer type. For instance, in the case of police patrols in New York City, "customers" refer to the incidents being reported, "customer type" refers to the priority the reported incident is assigned. Moreover, the occurrences of these incidents may follow a spatial Poisson distribution. Third, the requests for services from customers may arise in a deterministic or a random fashion or a combination of both. Finally, there may be some requirements regarding the nature of services that must be met. In all these applications of real-time CAD systems, two types of questions can be pursued: (1) strategic planning questions which concern problems such as the optimal fleet size, vehicles' capacities, etc., and (2) tactical planning questions which concern the optimal deployment of resources.

We note that, in every application area discussed earlier, the use of computers when they were first introduced was in the automation of manual operations, which had been performed reasonably well for several years. Moreover, we have, with some exceptions, ignored the use of the tremendous computational power a computer has in helping decision makers solve strategic or tactical planning problems. However, these exceptions, such as dial-a-ride services, vehicle routing and scheduling, and train dispatching show that the computational capabilities of the digital computer can be successfully employed to facilitate decision making on the real-time deployment of resources. In addition to these positive results, the development of computers from large, main frame computers to contemporary, less expensive, yet more powerful, local area mini computers has aroused our attention on this neglected, underutilized area which we shall call "real-time resource management."

Besides the monetary savings and other benefits generated from the real-time resource management in some of the application areas mentioned before, there are two other attractive advantages of real-time resource management which we have not yet pointed out. First, real-time resource management can avoid the "one-shot consultant" problem. Recall that real-time resource management involves utilizing the on-line computational capabilities of a digital computer to provide algorithmically derived support systems. Then, even if the consultant who developed such algorithmically derived support systems left the company, these support system, once built, can be operated by others and continue to aid in real-time decision-making. Second, real-time resource management creates job opportunities for less sophisticated people. One important goal in developing an algorithmically derived support system is to make it as "friendly" to users as possible. By "friendly," we mean easy to understand and operate. Once this goal is achieved, the second advantage of real-time resource management follows.

Finally, a difficulty associated with the implementation of a real-time CAD system should be addressed. Technological changes often involve behavioral changes and perceived and actual shifts of power within an organization, which, if overlooked, can become a real obstacle to the implementation of new technology.

## REFERENCES


1. Bell, W.J., L.M. Dalberto, M.L. Fisher, A.J. Greenfield, R. Jaikumar, P. Kedia, R.G. Mack, and P.J. Prutzman, "Improving the Distribution of Industrial Gases with an On-Line Computerized Routing and Scheduling Optimizer," **Interfaces**, 13(6):4-23, 1983.

2. Bodin, L., B. Golden, A. Assad and M. Ball, "Routing and Scheduling of Vehicles and Crews: State of the Art," **Computers and Operations Research**, 10(2):63-212, 1983.

3. Colton., K.W., M.L. Brandeau and J.M. Tien, "A National Assessment of Police Command, Control and Communications Systems," Public Systems Evaluation, Inc., Final Report 78-NI-AX-0144, August 1981.

4. Devanney, J.W., Offshore Oil Task Group, "The Georges Bank Petroleum Study," MIT Sea Grant Report MITSG-73-5, 1973.

5. Dietzel, G.F.L., A.W. Glass, P.J. Kleef, "SLIKTRACK-A computer simulation for the prediction of slick movement, removal of natural means, cleanup and potential damages arising from oil spills originating from offshore oil well-blow outs. Its development and application to the North Sea," Shell Report EP-47436, 1976.

6. Fisher, M.L., "The Lagrangian Relaxation Method for Solving Integer Programming Problems," **Management Science** 27(1):1-18, 1981.

7. Fisher, M.L., A. Greenfield, R. Jaikumar, "VERGIN: A Decision Support System for Vehicle Scheduling," Working Paper HBS 80-62, Division of Research, Harvard Business School, January 1982.

8. Fisher, M.L. and R. Jaikumar, "A Generalized Assignment Heuristic for Vehicle Routing," **Networks**, 11:109-124, 1981.

9. Glover, F. and D. Klingman, "Network Applications in Industry and Government," **AIIE Transactions**, 9(4):363-376, 1977.

10. Jaikumar, R, M.L. Fisher, A.J. Greenfield and P. Kedia, "Real-Time Scheduling of a Bulk Delivery Fleet: Practical Application of Lagrangian Relaxation," Working Paper HBS 83-84, Division of Research, Harvard Business School, 1983.

11. Larson, R.C., "Proactive Real-Time Management of Scarce Police Patrol Resources," MIT Operations Research Center proposal to National Institute of Justice.

12. Lin, S. and B. Kerningham, "An Effective Heuristic Algorithm for the Traveling Salesman Problem," **Operations Research**, 21:498-516, 1973.

13. Minkoff, A.S. and R.C. Larson, "Barge Operations System Simulator," Final Report No. 2P01141, ENFORTH Corporation, 1983.

14. Nyhart, J.D., H.N. Psaraftis and P.J. Yaroschak, "Putting an Oil Spill Cleanup Computer Model to Work for the Navy," **Naval Engineers Journal**, 1983.

15. Psaraftis, H.N., "A Dynamic Programming Solution to the Single Vehicle Many-to-Many Immediate Request Dial-A-Ride Problem," **Transportation Sciences**, 14(2), 1980.

16. Psaraftis, H.N, Water Pollution, McGraw-Hill 1982-83 Yearbook of Science and Technology.

17. Psaraftis, H.N, J.D. Nyhart, D.A. Betts, "First Experiences with the Massachusetts Institute of Technology Oil Spill Model," Proceedings of the Oil Spill Conference, San Antonio, Texas.

18. Psaraftis, H.N. and B.O. Zivgas, "A Tactical Decision Algorithm for the Optimal Dispatching of Oil Spill Cleanup Equipments," Working Paper OE-SP-82-1, Massachusetts Institute of Technology, 1983.

19. Sander, R.L. and W.M. Westerman, "Computer Aided Train Dispatching: Decision Support Through Optimization," **Interfaces**, 13(6):24-37, 1983.

20. Wilson, N.H.M., N.J. Colvin, "Computer Control of the Rochester Dial-A-Ride System," Massachusetts Institute of Technology, Center for Transportation Studies, Report Number 77-22, July 1977.

21. Wilson, N.H.M., R.W. Weissberg, and J. Hauser, "Advanced Dial-A-Ride Algorithms Research Project: Final Report," Massachusetts Institue of Technology, Department of Civil Engineering, Report R76-20, March 1976.

22. Wilson, N.H.M., E. Miller, "Advanced Dial-A-Ride Algorithms Research Project Phase II: Interim Report, Massachusetts Institute of Technology, Department of Civil Engineereing, Report R77-31, July 1977.