MANAGEMENT SCIENCE
Vol. 31, No. 12, December 1985
Printed in U.S.A.

A TACTICAL DECISION ALGORITHM FOR THE OPTIMAL DISPATCHING OF OIL SPILL CLEANUP EQUIPMENT*

HARILAOS N. PSARAFTIS AND BABIS O. ZIOGAS
Department of Ocean Engineering, Massachusetts Institute of Technology, Cambridge, Massachusetts 02139
IMI, Inc., New York, New York 10017

We develop an optimization procedure for assisting decision-makers in the allocation of resources for cleaning up a specific oil spill. The objective function is to minimize a weighted combination of spill-specific response and damage costs. Inputs to this problem include information about the outflow of oil, availability and performance of spill cleanup equipment, as well as costs of equipment transported and on-scene operation. A general (albeit separable) damage function is assumed. The algorithm is deterministic and is based on a dynamic program within which a series of 0-1 knapsack problems are solved repeatedly. Although this algorithm is approximate, its worst-case performance is quantified and we argue that under realistic inputs the procedure can be expected to produce solutions very close to optimality. Under prescribed conditions we prove that the algorithm produces optimal solutions. A realistic example based on the *Argo Merchant* oil spill is presented to provide insight into the structure of this problem. Finally, we discuss possible uses of this model within the existing and alternative operational and policy environments.
(DYNAMIC PROGRAMMING—APPLICATIONS; ENVIRONMENTAL MANAGEMENT)

1. Introduction

Massive catastrophic oil spills such as the *Torrey Canyon* spill in Britain (1967), the *Amoco Cadiz* spill in France (1978), and the IXTOC-1 spill in Mexico (1979) always focus public attention on the damage caused by such accidents. While such massive spills are very rare, many more small and moderate spills occur on a daily basis from a variety of sources: operational discharges from tankers, vessel collisions, pipeline ruptures, etc. As a result, oil spills have become an everyday concern for governments, as well as for the ocean transportation, oil exploration and oil production industries, and a whole spectrum of actions are taken or considered in order to alleviate oil spill problems. A significant part of such actions falls into the *response* category.

Oil spill response concerns the emergency action that must be taken so that pollution of the sea and coastline is kept under control once an oil spill occurs. Part of such action involves the dispatching of specialized cleanup equipment to the spill site in order to contain and recover the spilled oil. Equipment may be a combination of booms, which are protective barriers that help prevent the uncontrollable spreading of the oil; skimmers, which are devices that pump and recover the oil; barges, which can be used to transport the recovered oil to a disposal site; sorbents, which are materials to absorb the oil, and a variety of other means. In addition, chemicals can be used in order to disperse the slick. The real-world decision-maker (in the case of the Coast Guard, the On Scene Coordinator) is faced with an extremely complex problem, having to address such issues as availability of equipment, performance degradation with bad weather, and uncertain movement of the oil slick, and having also to balance the potentially high and uncertain cost of oil spill damage with the similarity high cost

*Accepted by Warren E. Walker; received January 4, 1983. This paper has been with the authors 7 months for 2 revisions.

0025-1909/85/3112/1475$01.25

of cleanup operations. The problems become even more complex for the strategic planner who must recommend "optimal" stockpiles of cleanup equipment so as to respond to spills that may occur in the future.

Although the core of the oil spill cleanup problem seems to be the question of the evaluation of the tradeoffs between cleanup and damage costs, literature has typically ignored those tradeoffs. Most of the literature to date has dealt with laboratory and theoretical research on the behavior and effects of oil at sea, with specific case studies that assess the environmental or economic impact of a specific spill, or with work on the performance of oil spill recovery equipment. Very few studies to date have addressed the decision-making process in oil spill cleanup (see, for instance, Fraser and Cochran 1975, Dietzel et al. 1976, Conrad 1978, Charnes et al. 1979, TSC 1979, Seaward International 1979, Versar 1981, and Belardo et al. 1984). However, while all of these studies have merit, they are not enough to make the state-of-the-art in oil spill decision-making comparable to the state of knowledge in other emergency service environments, such as the urban one, for which there is a vast literature covering both methodology and applications (see, for instance, Walker et al. 1979, among many others).

The purpose of this paper is to describe an analytical methodology developed to assist the decision-maker to make an optimal allocation of resources for cleaning up *a specific spill* after its occurrence is made known. This work is part of a broader Massachusetts Institute of Technology project that began in July 1979 with support from a consortium of government and industry organizations. The goal of the overall project has been to create a computerized tool that would provide the user with the ability to analyze complex decisions regarding oil spill cleanup. The model—whose generic structure is shown in Figure 1—integrates all relevant parts of a spill response system and explicitly incorporates analytical descriptors of system performance as well as decision-making techniques. The description of the entire model structure is beyond the scope of this paper—the interested reader is referred to Psaraftis (1982), Psaraftis et al. (1983, to appear), Ziogas (1982), and Tharakan (1982) for more details. This paper focuses on the so-called "Tactical-Model" (see Figure 1).

THE MIT OIL SPILL MODEL

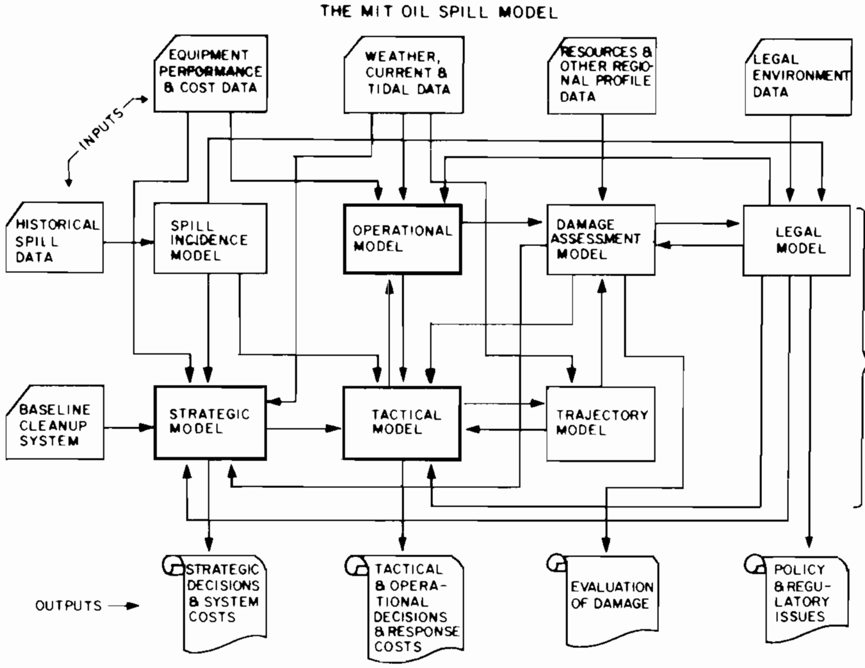

GURE 1. Structure of the MIT Oil Spill Model.

Drawing from Anthony's hierarchical framework for the analysis of business systems (Anthony 1965), one can divide the oil spill decision-making process into three hierarchical levels: *strategic*, *tactical*, and *operational*. Various definitions of these levels appear in the literature. In this paper, we use the following definitions:

(1) The *strategic* level, where one wishes to determine the quantities, types and locations of equipment that should be stockpiled to respond to potential future oil spills;

(2) The *tactical* level, where one wishes to determine aggregate actions that should be taken to respond to a specific spill, such as what equipment should be dispatched on scene, how long that equipment should stay on scene, etc.; and

(3) The *operational* level, where one examines in much more detail actions that must be taken on scene, such as geometric configuration of boom deployment, skimming, dispersants, etc., or spatial allocation of cleanup resources to protect sensitive areas.

This paper develops algorithms for the *tactical* decision level. It assumes that strategic variables have already been decided upon, and that operational variables will be decided upon after the tactical-level decisions are made. Of course, decisions made at the upper levels are constraints on the problems defined at the lower levels. Conversely, information from the performance of the system at the lower levels may be used to feed back into the problems defined at the upper levels.

It is also clear that although the line between the strategic and tactical decision levels is explicitly defined, the distinction between tactical and operational decisions is arbitrary (for this problem). Both tactical and operational decisions concern actions in response to a specific spill. However, complexity considerations suggest that it would be unwise to include simultaneously all of such actions (from the most aggregate to the most detailed) into a single decision problem. The distinction between the two levels is made so as to both capture the aggregate features of the response (tactical level), and allow greater user discretion on the more detailed actions on scene (operational level). The operational-level actions are much more unstructured or even impossible to model.

The remainder of the paper is organized as follows:

In §2 we define a generic version of the tactical decision problem and formulate it as an optimization problem. §3 develops both an exact and an approximate solution methodology, both based on Dynamic Programming, for solving the above problem. It is seen that the approximate procedure is substantially more tractable than the exact one and, under realistic assumptions, produces near optimal or optimal solutions. §4 presents a realistic illustration of the model in conjunction with the *Argo Merchant* oil spill. The illustration highlights possible uses of the model within the existing and alternative operational and policy environments. Those uses are further discussed in §5, together with other issues and suggestions for further research.

2. Definition of the Tactical Decision Problem

We start the description of the tactical decision problem with the drastic assumption that the problem is deterministic, even though the inputs to this problem include variables such as spill outflow rate or weather conditions, which are inherently unpredictable in the real world. However, we must learn to walk before we can run. As it will be seen below, the tactical decision problem has nontrivial difficulty even in its deterministic version. Moreover, the solution to the deterministic version will be seen to provide significant insights into the structure of the problem, insights that may be useful in a stochastic extension. Such an extension has not been examined in this paper. The description of a generic "tactical" oil spill scenario follows:

At some specific time and at some geographical location, an oil spill of known oil type and quantity occurs. The spill is reported to the responsible authorities at a given notification time after it occurs. A variety of oil spill cleanup "equipment sets" are located at known sites and are available for dispatch to the spill site. The performance of each equipment set is a known function of the prevailing weather conditions. Future weather conditions and spill outflow are known with certainty. The problem can be phrased as follows: If the costs of using each equipment set in cleaning up the spill are known, and if the damage costs that would occur if the oil impacts sensitive areas are known, what is the "response tactic" that minimizes a weighted combination of spill-specific cleanup and damage costs?

Some important clarifications follow:

(1) By oil spill cleanup "equipment set" we mean an integrated, sufficiently equipped and self-contained *package* capable of "cleaning up" a prescribed quantity of oil in a given period. "Cleaning up" is interpreted here in a generic sense to mean anything that effectively prevents the oil from impacting sensitive areas. It can mean that the oil is "contained and removed by mechanical means and then stored, transported and disposed of;" that the oil is "treated and dispersed by chemical means into the water column;" or that some other effective measure is taken. An equipment set is composed of several individual components. For instance, an integrated skimming system is typically a combination of booms, pumps and storage capacity. The dispatching time, cost, and performance of the set depend on the dispatching times, costs and performances of all its separate components (which need not be stockpiled all at the same location). In this paper we shall not attempt to analyze how the performance of each part affects the whole. Rather, we assume that we are dealing only with a specified number of *sets*, all of which have known performances. Similarly we shall not be concerned here with the mechanisms that make the above performance dependent on the prevailing weather conditions, but rather assume that the performance is a known input (for details on the modeling of equipment performance see Ziogas 1982). The rationale of using equipment *sets* instead of individual components, as well as how one goes about defining those sets is further discussed in §5.2.

(2) All costs of using cleanup equipment are assumed to be spill-specific, variable (opportunity) costs. In other words, equipment acquisition costs and other costs that have been committed at the strategic decision level are *sunk costs* and are not considered in the tactical problem.

(3) An oil spill typically spreads onto a surface area of considerable size, and consists of individual components called *spillets*. The behavior of each spillet is tracked by the "Trajectory" part of the overall model as a function of prevailing weather conditions and the characteristics of the spill. In our model, the response at the tactical level is only an *aggregate response*. It does not attempt to "optimize" the spatial distribution of cleanup resources among spillets. In that respect, the tactical level assumes that all equipment is distributed to all spillets *uniformly*, namely, that a spillet of 3,000 gallons of oil would receive, on the average, three times more cleanup capability than a spillet of 1,000 gallons of oil, irrespective of the location of those spillets. The tactical level also assumes that equipment mobilization can only control the *volume* and not the trajectory of each spillet. Those assumptions of course concern only the tactical level of the problem. Once an aggregate response is established at that level, the user can enter the operational level where he can make further decisions on how to allocate that response across the geographical area of interest. Implicit in all the above is the assumption that potential damages accounted for in the evaluation of tactical decisions *do not* depend on decisions at the operational level, but only on tactical response decisions. A discussion of what happens if we relax this assumption (that is, if we allow such feedback from the operational level into the tactical level) is presented in §5.3. The way damages are calculated is explained later on in this section.

(4) Decisions at the tactical level can be further broken down into two levels of aggregation: At the more aggregate level are decisions on the "aggregate cleanup capability" necessary to respond to the spill, expressed in volume of oil cleaned up per unit of time and throughout the duration of the spill. At the more detailed level are decisions on what specific equipment sets should be dispatched on scene, including from what locations the dispatching should originate, throughout the duration of the spill. Decisions at an even more detailed level (such as exactly how should equipment be spatially deployed, how should it be operated, etc.) belong to the operational level of the problem and will not be considered here. We assume that the decision-maker accounts for all events and evaluates decisions at *discrete points in time.* These points, called *time stages*, are separated by a user-specified time interval (e.g. 3 hours, 6 hours, etc.). A sequence of such decisions is what is termed a "response tactic" for the spill in question. It is clear that a response tactic is uniquely defined once the sequence of decisions at the more detailed level is known, for that sequence of decisions is known, the aggregate cleanup capability is known as well. The main justification for the consideration of the response tactic at two levels of aggregation is computational efficiency, as we shall explain further in §3.

(5) The objective for the tactical decision problem is to minimize the sum of two cost components: the spill-specific cleanup costs (assumed to be weighted by 1), and the damage costs, multiplied by a *user-specified* weight. By systematically varying the weight of the damage costs, the user can adjust the relative emphasis of one of the two cost components versus the other and evaluate the tradeoffs between them. More discussion on this point is presented in §5.4.

Based on the above we are in a position to define the input variables to the tactical decision problem as follows:

1. W: The user-specified weight by which damage costs are multiplied. Cleanup costs are assumed to be weighted by 1.0 ($0 \leqslant W < +\infty$).

2. Dt: The user-specified time interval between two consecutive time stages (in hours).

3. NOTFI: The time stage corresponding to the spill notification time.

4. NOUTFL: The time stage corresponding to the "end of oil spill outflow," beyond which there is no more discharge of oil into the sea.

5. NCLUP: An upper bound for the time stage corresponding to the "end of cleanup operations," beyond which all equipment sets return to their bases.

6. NENDSP: An upper bound for the time stage corresponding to the "end of the spill event", beyond which no more events related to the spill in question are significant enough to be accounted for.

7. d_n ($n = 1, \ldots,$ NOUTFL): The volume of oil that is discharged during the time interval $[nDt, (n+1)Dt]$. This volume is assumed to be discharged instantaneously at time stage n (in gallons).

8. NOEQ: Number of distinct equipment sets available.

9. t_{in} ($i = 1, \ldots,$ NOEQ; $n = 1, \ldots,$ NCLUP): The time (in time stages) it takes equipment set i to arrive on scene if it arrives there at time stage n. t_{in} consists of mobilization delays plus the transport and set-up times needed to have *all* components of set i on scene and operable. t_{in} depends on n since the "spill location" (defined for the tactical decision problem as the location of the centroid of the spill) will generally move due to wind, currents, etc.

10. r_{in} ($i = 1, \ldots,$ NOEQ; $n = 1, \ldots,$ NCLUP): The volume of oil that can be "cleaned up" if equipment set i operates on scene between time stages n and $n + 1$ (in gallons). r_{in} is time-dependent since the performance of set i depends on weather conditions which may change with n.

11. fc_i ($i = 1, \ldots$, NOEQ): Fixed costs incurred whenever equipment set i is dispatched on scene. This includes costs such as mobilization, transportation and set-up, but does not depend on how long set i stays on scene.

12. c_{in} ($i = 1, \ldots$, NOEQ; $n = 1, \ldots$, NCLUP): Operating cost of equipment set i if it spends time on scene between stages n and $n + 1$.

Some of the above inputs (such as t_{in}, r_{in}, c_{in}) are *derived* inputs, that is, depend on inputs such as weather conditions, manufacturer's specifications, spill location etc., which are not explicit inputs to the tactical problem. Those inputs are furnished to the tactical model via other parts of the overall model (Trajectory, Operational etc.).

Decision variables for the tactical decision problem are the following ($n = 1, \ldots$, NENDSP; $i = 1, \ldots$, NOEQ):

1. X_n: Aggregate cleanup capability on scene at time stage n. This is the amount of oil "cleaned up" between time stages n and $n + 1$.

2. $Y_{in} = 1$ if equipment set i is on scene at time stage n, 0 otherwise.

Optimal tactical decisions depend not only on the cleanup cost information mentioned above, but also, and with equal importance, on damage information. Assume that the volume of oil at sea at time stage n is equal to A_n. A certain amount of oil will escape (become nonrecoverable) between n and $n + 1$. The "escaping" process is highly complex and nonlinear. It is governed by phenomena such as evaporation, drifting, natural dispersion etc., and also depends on the cleanup capability on scene at the time.

In this paper we shall assume that the volume of oil escaping cleanup between time stages n and $n + 1$ is furnished externally by the "escape function" $f_n(A_n, d_n, X_n)$. f_n is furnished by the "Trajectory Model" (Figure 1) which takes into account information such as winds, currents, water depth, type of oil, etc., and returns outputs such as amount of oil evaporated, dispersed, and emulsified, tracks down the movement of the oil in the area of interest, and determines parameters such as thickness and surface area of the slick (see Ziogas 1982).

The portion of the nonrecoverable oil that will impact environmental and economic resources in the area will create damages. In this paper we shall assume that the damage costs due to the escaped oil f_n are furnished externally by a function $D_n(A_n, d_n, X_n, f_n)$ (damage function). Like f_n, D_n is in itself part of a complex algorithm that takes into account information about toxicity of oil, location of resources in the area of interest, and includes the value of the lost oil (Damage Assessment Model of Figure 1). The reader should refer to Demis (1984) for a detailed discussion of the methodology used to construct such a function. D_n is usually highly nonlinear due to "threshold" effects that can take place when oil hits the shoreline. However in this paper we assume that damages are time-wise *separable*: they can be broken down by time stages. A discussion of the separability assumption and other aspects of the damage function is presented in §5.1.

Based on the above we can write the problem's objective function as follows:

$$\text{Minimize} \quad \sum_{n=1}^{\text{NENDSP}} \sum_{i=1}^{\text{NOEQ}} \left[c_{in} Y_{in} + fc_i Y_{in}(1 - Y_{i,n-1}) \right] + W \cdot \sum_{n=1}^{\text{NENDSP}} D_n(A_n, d_n, X_n, f_n). \tag{1}$$

There are a number of constraints in our problem:

(1) Clearly, the aggregate cleanup capability X_n is related to the vector of Y_{in}'s:

$$X_n = \sum_{i=1}^{\text{NOEQ}} r_{in} Y_{in}, \qquad n = 1, \ldots, \text{NENDSP}. \tag{2}$$

(2) An equipment set cannot be on scene earlier than the time it takes to notify, mobilize, dispatch, and set-up that set. Hence

$$Y_{in}(t_{in} + \text{NOTFI} - n) \leqslant 0, \quad i = 1, \dots, \text{NOEQ}; \quad n = 1, \dots, \text{NENDSP}. \quad (3)$$

(3) There is conservation of mass for the oil between n and $n + 1$. The transition equation is given by:

$$A_{n+1} = \text{Max}[0, A_n + d_n - f_n - X_n], \quad n = 1, \text{NENDSP} - 1, \quad (4)$$

where the maximum is taken since it is conceivable that $A_n + d_n - f_n - X_n < 0$ (excess capability on scene).

(4) Finally, for $n = 1, \dots, \text{NENDSP}$; $i = 1, \dots, \text{NOEQ}$:

$$X_n \geqslant 0; \quad Y_{in} = 0 \text{ or } 1; \quad Y_{i0} = 0; \quad A_n \geqslant 0. \quad (5)$$

3. Solution of the Problem

This section is structured in three parts. First, we present an *exact* Dynamic Programming algorithm for solving the tactical decision problem outlined in the previous section. For anything but very small problem sizes such an approach is intractable from a running time and storage requirement viewpoint. Second, we present an *approximate* version of the above algorithm for solving the same problem. The approximate version can solve the problem with only a fraction of the running time and storage required by the exact version. Third, we derive bounds that quantify the approximate algorithm's *worst-case* performance (in terms of deviation from the exact optimum) and argue that the algorithm's *average* performance can be expected to be very good under realistic input values and optimal under special circumstances.

3.1. *Exact Algorithm*

Define $V_n^*(A_n, \bar{Y}_{n-1})$ as the minimum achievable total weighted cost from time stage n until the end of the event, given that at stage n the "state" of the event is as follows:

(a) Volume of oil at sea: A_n.

(b) Cleanup equipment on scene between stages $n - 1$ and n described by the vector $\bar{Y}_{n-1} = (Y_{1,n-1}, \dots, Y_{\text{NOEQ},n-1})$.

Decision variable for the next stage is $\bar{Y}_n$, the vector defining what equipment should be on scene from n to $n + 1$. It is then clear that $V_n^*(A_n, \bar{Y}_{n-1})$ obeys the following recursive formula:

For $1 \leqslant n < \text{NCLUP}$:

$$V_n^*(A_n, \bar{Y}_{n-1}) = \underset{\bar{Y}_n \in E_n}{\text{Min}} \Bigg[\sum_{i \in I_n} c_{in} Y_{in} + \sum_{i \in I_n} fc_i Y_{in}(1 - Y_{i,n-1}) + W \cdot D_n\Big(A_n, d_n, \sum_{i \in I_n} r_{in} Y_{in}, f_n\Big) + V_{n+1}^*(A_{n+1}, \bar{Y}_n)\Bigg]. \quad (6)$$

For $n = \text{NCLUP}$:

$$V_n^*(A_n, \bar{Y}_{n-1}) = W \cdot \sum_{k=1}^{\text{NENDSP}} D_k(A_k, d_k, 0, f_k), \quad (7)$$

where, for $1 \leqslant n < \text{NENDSP}$:

$$I_n = \begin{cases} \{i : t_{in} + \text{NOTFI} - n \leqslant 0\} & \text{if } 1 \leqslant n < \text{NCLUP}, \\ \varnothing & \text{otherwise}; \end{cases} \quad (8)$$

$$E_n = \{\bar{Y}_n : Y_{in} = 0 \text{ or } 1 \text{ for } i \in I_n, 0 \text{ otherwise}\}; \quad (9)$$

$$A_{n+1} = \text{Max}\Big[0, A_n + d_n - \sum_{i \in I_n} r_{in} Y_{in} - f_n\Big]. \quad (10)$$

In (7) damages accrue according to a "do-nothing" (or, euphemistically, "surveillance and monitoring") tactic for $n \geqslant$ NCLUP. (8) defines the equipment sets that can physically be on scene at stage n, while (9) defines the feasible combinations of vector $\overline{Y}_n$. Relation (10) is the equivalent of (4) and defines A_{n+1} in both (6) and (7). Finally, in (6), (7) and (10) we have written f_n instead of $f_n(A_n, d_n, \sum_{i \in I_n} r_{in} Y_{in})$ for convenience purposes. Of course, all terms of the form $\sum_{i \in I_n} r_{in} Y_{in}$ vanish for $n \geqslant$ NCLUP.

The optimal value of the tactical decision problem is $V_1^*(0, \overline{0})$ where $\overline{0}$ is the zero vector for the value of $\overline{Y}_0$.

It is clear that the storage requirement associated with this exact algorithm is $O(\text{NVOL} \cdot \text{NCLUP} \cdot 2^{\text{NOEQ}})$ where NVOL is the number of values necessary to discretize A_n. In addition, the fact that we have a maximum of 2^{NOEQ} alternatives for equipment dispatching at the next stage makes the total computational effort of this algorithm grow as $O(\text{NVOL} \cdot \text{NCLUP} \cdot 2^{2\text{NOEQ}})$, a clearly intractable function for anything but trivial problems.

3.2. *Approximate Algorithm*

The algorithm developed above is intractable mainly because we have to keep in the state space vector $\overline{Y}_{n-1}$, whose domain has size that is an exponential function of the number of equipment sets. Suppose now that instead of $\overline{Y}_{n-1}$ we decide to use X_{n-1}, the aggregate on scene capability between $n-1$ and n, as a *surrogate* representation of equipment deployment on scene. Clearly, specifying X_{n-1} instead of $\overline{Y}_{n-1}$ gives us only a partial idea of the equipment configuration between $n-1$ and n, for there may be a great number of ways equipment sets can be combined to produce a total cleanup capability of X_{n-1}.

The approximate approach assumes that the only possible combination of equipment sets corresponding to a given aggregate capability is the one that guarantees at least such an aggregate capability *while minimizing total operating costs during the interval.* In other words, given X_{n-1}, the approximate approach will assume that equipment vector $\overline{Y}'_{n-1}$ solves the following 0-1 knapsack problem:

$$\begin{aligned} \text{Min} \quad & \sum_{i \in I_{n-1}} c_{i,n-1} Y'_{i,n-1} \\ \text{st} \quad & \sum_{i \in I_{n-1}} r_{i,n-1} Y'_{i,n-1} \geqslant X_{n-1}, \\ & Y'_{i,n-1} = 0 \text{ or } 1; \qquad i \in I_{n-1}. \end{aligned} \tag{11}$$

Decision variables for the next stage will similarly be X_n instead of $\overline{Y}_n$ and again we shall assume that such capability is provided by equipment whose vector $\overline{Y}'_n$ solves the following 0-1 knapsack problem:

$$\begin{aligned} \text{Min} \quad & \sum_{i \in I_n} c_{in} Y'_{in} \\ \text{st} \quad & \sum_{i \in I_n} r_{in} Y'_{in} \geqslant X_n, \\ & Y'_{in} = 0 \text{ or } 1; \qquad i \in I_n. \end{aligned} \tag{12}$$

With state vector (A_n, X_{n-1}), we then define $V_n(A_n, X_{n-1})$ as the "approximate optimal value function," which *by definition* obeys the following recursive relationship:

For $1 \leqslant n <$ NCLUP:

$$\begin{aligned} V_n(A_n, X_{n-1}) = \underset{X_n}{\text{Min}} \Bigg[& \sum_{i \in I_n} c_{in} Y'_{in} + \sum_{i \in I_n} fc_i Y'_{in}(1 - Y'_{i,n-1}) \\ & + W \cdot D_n\Big(A_n, d_n, \sum_{i \in I_n} r_{in} Y'_{in}, f'_n\Big) + V_{n+1}(A'_{n+1}, X_n) \Bigg]. \end{aligned} \tag{13}$$

For n = NCLUP:

$$V_n(A_n, X_{n-1}) = W \cdot \sum_{k=n}^{\text{NENDSP}} D_k(A_k, d_k, 0, f'_k), \tag{14}$$

where I_n is defined by (8), $\overline{Y}'_{n-1}$ and $\overline{Y}'_n$ by solving (11) and (12) respectively, and where for $1 \leqslant n <$ NENDSP:

$$A'_{n+1} = \text{Max}\left[0, A_n + d_n - \sum_{i \in I_n} r_{in} Y'_{in} - f'_n\right]. \tag{15}$$

In (13), (14) and (15), we have written f'_n instead of $f_n(A_n, d_n, \sum_{i \in I_n} r_{in} Y'_{in})$ for convenience purposes. Again, the term $\sum_{i \in I_n} r_{in} Y'_{in}$ vanishes for $n \geqslant$ NCLUP.

$V_n(A_n, X_{n-1})$ is only *an approximation* of the minimum total weighted cost from stage n until the end of the event for the following reasons: First, the equipment configuration on scene at stages $n-1$ and n is restricted to the one that solves (11) and (12) respectively. Second, it may happen that the combination of equipment sets that solves (11) and (12) has an actual aggregate capability *strictly greater* than X_{n-1} or X_n respectively. Such a possibility is partially ignored in the recursion, which always considers X_{n-1} and X_n as state variables (as opposed to $\sum_{i \in I_{n-1}} r_{i,n-1} Y_{i,n-1}$ and $\sum_{i \in I_n} r_{in} Y_{in}$ respectively). Computationally, this approach is much more tractable than the previous one, for it replaces 2^{NOEQ} states with the number of states necessary to discretize X_{n-1}.

Next we investigate how close the approximate algorithm can approach the optimum produced by the exact algorithm.

3.3. *Worst-Case and Average Performance of the Approximate Algorithm*

To find out how much the solutions of the approximate algorithm deviate from the optimal solutions produced by the exact algorithm, we derive a set of relationships between $V_n(A_n, X_{n-1})$ and $V_n^*(A_n, \overline{Y}_{n-1})$. The only assumptions with regard to the "realism" of f_n and D_n are the following: For all n (a) both f_n and D_n are nondecreasing functions of A_n and nonincreasing functions of X_n; (b) for any $\Delta A_n \neq 0$ the ratio of $\Delta f_n / \Delta A_n$ is between 0 and 1 (everything else being equal), and finally (c) D_n is a nondecreasing function of f_n. Under those assumptions, the relationships between $V_n(A_n, X_{n-1})$ and $V_n^*(A_n, Y_{n-1})$ can be summarized as follows:

(1) For n = NCLUP, $V_n(A_n, X_{n-1}) = V_n^*(A_n, \overline{Y}_{n-1})$ for any A_n, X_{n-1} and $\overline{Y}_{n-1}$.

(2) For $1 \leqslant n <$ NCLUP,

$$V_n^*\left(A_n, \overline{Y}_{n-1}\right) \leqslant V_n(A_n, X_{n-1}) \leqslant V_n^*\left(A_n, \overline{Y}'_{n-1}\right) + \sum_{k=n}^{\text{NCLUP}-1} \sum_{i \in I_k} fc_i,$$

for any A_n, and for any $\overline{Y}_{n-1}$, X_{n-1} and $\overline{Y}'_{n-1}$ such that

(a) $X_{n-1} = \sum_{i \in I_n} r_{i,n-1} Y_{i,n-1}$;

(b) $\overline{Y}'_{n-1}$ solves the knapsack problem associated with X_{n-1} (11).

(3) As a corollary of (2),

$$V_1^*(0, \overline{0}) \leqslant V_1(0,0) \leqslant V_1^*(0, \overline{0}) + \sum_{k=1}^{\text{NCLUP}-1} \sum_{i \in I_k} fc_i.$$

The above relationships can be proven by induction (the proofs are omitted from this paper due to space limitations but are available from the authors). The implications of (3) are twofold:

(i) The value produced by the approximate algorithm is bounded from below (as expected) by the optimal value of the problem and cannot exceed that value by more than the quantity $\sum_{k=1}^{\text{NCLUP}-1} \sum_{i \in I_k} fc_i$.

(ii) The approximate algorithm produces *optimal* solutions if all fixed dispatching costs are zero. That is, if $fc_i = 0$ for $i = 1, \ldots,$ NOEQ, we incur no error by substituting the detailed equipment information $\overline{Y}_n$ with the more aggregate equipment information X_n.

Not unexpectedly, we can see that the approximate algorithm cannot always guarantee an optimal solution in case some or all fc_i's are nonzero. This is, in a sense, the "price" one has to pay in order to be able to solve this problem more efficiently than by using the exact approach. We should however emphasize that the upper bound of the absolute error of the approximate algorithm $(\sum_{k=1}^{\mathrm{NCLUP}-1} \sum_{i \in I_k} fc_i)$ is only a worst-case figure, and a rather loose one for that matter. Indeed, such a worst-case error can be realized if and only if (a) the approximate algorithm dispatches every equipment set back and forth at every time stage of the spill *and* (b) the exact algorithm recommends no response to that spill. Not only were we unable to concoct a pathological case in which such a (maximum) error is realized, but more importantly we proved that irrespective of the values of the inputs and the forms of f_n and D_n, it is *impossible* for the absolute error to be *exactly* equal to the above upper bound. This means of course that the upper bound is not tight, and that it is conceivable that a lower worst-case error could be derived after a more involved analysis. Such an analysis is, of course, outside the scope of this paper.

The above considerations refer to an investigation of the approximate algorithm from a *worst-case* viewpoint. However, such worst-case analyses are of theoretical interest only because the pathological cases in which such worst-case performance is realized are usually very unlikely to occur. What is likely to be of more interest to the algorithm's user is its performance in typical problem cases "in practice" or "on the average". In our case, and from an *average* performance viewpoint, we feel that the approximate algorithm is likely to perform very well, for the following reasons: (a) If weather conditions (and hence, equipment performance) do not fluctuate dramatically over time, dispatching equipment back and forth several times is a rather unlikely outcome; (b) In those cases in which multiple equipment dispatching is recommended by the approximate algorithm because of (say) dramatic changes in weather conditions, such multiple equipment dispatching is likely to be recommended by the exact algorithm as well. The two algorithms handle fixed dispatching costs in essentially the same fashion (compare equations (6) and (13)) and are likely to recommend similar response tactics even under drastic fluctuations in weather conditions over time. Finally, (c) in many actual situations, equipment utilization is billed by the hour or by the day. Dispatching costs are thus included as part of the operating costs, making fixed dispatching costs much lower than total operating costs or total weighted costs. All of these reasons argue in favor of the approximate algorithm's performance.

Based on the above, we decided to use the approximate DP algorithm for solving the tactical decision problem. To further enhance its computational tractability, and despite the fact that this was not necessary, the 0-1 knapsack problems within the recursion were solved heuristically, by using the following "greedy" heuristic for a problem of knapsack size X_n (12):

Step 1. Rank-order all equipment sets $i \in I_n$ by nondecreasing order of c_{in}/r_{in} ratios.

Step 2. Following the above established order, choose as many sets as necessary for X_n to be satisfied. For those sets, put $Y_{in} = 1$. For all other sets, put $Y_{in} = 0$.

The greedy heuristic runs in $O(|I_n| \log |I_n|)$ time (by contrast to an exact approach that runs in $O(|I_n| X_n)$ time). This heuristic's *worst-case* error ratio is equal to

$$\frac{\sum_{i \in I_n} c_{in}}{X_n} \cdot \operatorname*{Max}_{i \in I_n} \left(\frac{r_{in}}{c_{in}} \right).$$

That is, the greedy heuristic is expected to perform better in cases X_n is high and/or the maximum r_{in}/c_{in} ratio is low. From an *average* performance viewpoint the greedy heuristic is expected to perform very well and it also provides a convenient rule of thumb in selecting equipment on a cost/benefit basis.

The approximate approach proved viable for problems involving up to 40 different equipment sets and about 50 time stages. We feel that this can be improved by further code refinements. Under realistic input values we never observed any symptoms associated with a worst-case performance (such as sending equipment back and forth many times). The code was written in FORTRAN and implemented on a VAX 11/782 at MIT.

4. An Illustrative Application

By far the most important illustrative application of the entire MIT Oil Spill Model has been within the New England regional context. The application necessitated a comprehensive data collection effort to compile information on all parts of the input (Figure 1), including data on environmental and economic resources in the area, geomorphology of the shoreline, cleanup equipment inventory, historical spill data, etc. In this section we shall present a flavor of the *tactical* part of this application.

The tactical application has been an "after-the-fact" analysis of actions, events, costs and damages connected with the *Argo Merchant* oil spill. On December 15, 1976, the tanker *Argo Merchant*, carrying 7.7 million gallons of heavy crude, ran aground 27 nautical miles southeast of Nantucket Island, Massachusetts. Small quantities of oil started leaking from the grounded tanker, gradually increasing during the next four days. On December 20, two million gallons escaped from the vessel and, on December 21, the ship broke in two, releasing about three million gallons. The next day the bow section broke again, releasing the remaining oil.

An extensive mobilization of cleanup equipment, manpower and scientific support was orchestrated by the U.S. Coast Guard. Fortunately, prevailing northwesterly winds made the oil drift offshore. Although the mobilization was massive, no oil was offloaded from the tanker or removed from the sea. After years of research and litigation, the only documented damage has been the value of the lost oil (estimated at between $4–$5 million). Although no oil was ever recovered, the total cost of cleanup equipment mobilization was also high: White and Nichols (1983) estimated the total cleanup cost of that spill at about $1.8 million (although it is not clear how they arrived at that figure).

We choose the *Argo Merchant* spill to test the model for several reasons: First, we considered it important to run the model with input data that were reasonably well documented; and, indeed, the literature covering that spill is rich. Second, what stimulated our interest was not so much an account of the actual event, but the opportunity we saw in using the model to address important "what if" questions regarding that spill. We asked the following kinds of questions: What would have happened if winds in the *Argo Merchant* oil spill were blowing in the opposite direction? What if the spill had occurred in the midst of the summer tourist season and winds carried it on to the New England coast? What if certain categories of cleanup equipment could or could not be used during the incident? We analyzed the above and other related questions from the point of view of oil spill trajectory, equipment dispatching, cleanup costs and damage assessment.

Information on oil outflow and weather conditions was based on many sources, the two basic ones being Milgram (1977) and Pollack and Stolzenbach (1978). The equipment database consisted of 40 different integrated equipment sets. All major equipment actually mobilized or standing by during the incident (such as ADAPTS systems, etc.) was included. This database was broken down into four major cleanup techniques: offloading, mechanical removal with storage, mechanical removal without storage, and chemical dispersants. The user of the model had the option to choose *one* of the above four techniques or any combination of them. Performance and cost data of the equipment were based on best estimates from experimental results, manufacturer specifications and contractor data sheets. A separate subroutine within the Operational Model was used to calculate the performance of each equipment set as a function of the performance of its individual components (pumps, storage, etc.), of prevailing weather conditions, and of oil type. The database included dispatching time estimates for each of the sets to arrive at the incident site. Damage assessment data were based on extensive information available on fisheries, tourism, property values, etc., for the region of interest. The offloading option was examined in *some* of the runs despite the bad weather conditions that prevailed at times (see Ziogas 1982 for complete details on the above inputs and the assumptions behind them).

Several runs related to the *Argo Merchant* case were performed. In this paper we focus on two major categories: First, runs related to the *actual* incident, and, second, runs related to a *worst-case* scenario of that incident. The worst-case variant differed from the actual case only in three respects: First, wind direction was shifted by 190° (clockwise), so that the spill could move toward the Massachusetts coast. Second, occurrence of the incident was changed from winter to summer, so that the tourist industry would be at its peak. And third, the oil was assumed to be light Diesel No. 2 instead of crude, so that its toxicity would cause more damage.

Several kinds of response were considered for both the actual and worst-case scenarios. In this paper we present results on the following response options:

(1) The "do-nothing" response, in which the model was forced to "benignly neglect" the spill (in both actual and worst-case scenarios). Although a politically unacceptable option, this was in fact the only model option that matched the actual event, since no oil was recovered. This could also establish an upper bound on the level of damages.

(2) The "optimal" response, in which the model minimized the sum of cleanup plus damage costs, with *any* combination of cleanup techniques allowed (in both actual and worst-case scenarios).

(3) Same as (2) but with only mechanical removal equipment (with storage) allowed (in worst-case scenario only).

(4) Same as (2) but with only mechanical removal equipment (without storage) allowed (in worst-case scenario only).

(5) Same as (2) but with only chemical dispersants allowed (in worst-case scenario only).

Table 1 summarizes cleanup and damage costs for options (1) and (2) above. From a "trajectory" viewpoint, the model accurately predicted the movement of the slick in the "actual" scenario/"do-nothing" option. Figure 2 is a computer graphics output of the "worst-case" scenario/"do nothing" option, showing the slick impacting the Massachusetts and Rhode Island shoreline resource grid system.

Several comments can be made regarding Table 1:

TABLE 1

Summary of Cleanup and Damage Costs for Two Scenarios and Two Response Options. For Further Explanations See Text.

Scenario	"Actual"		"Worst-Case"	
Response Option	(1)	(2)	(1)	(2)
Cleanup Cost	—	0.086^M	—	0.034^M
Damage Cost	4.6^M	0.79^M	34.8^M	1.1^M
Total Cost	4.6^M	0.876^M	34.8^M	1.134^M

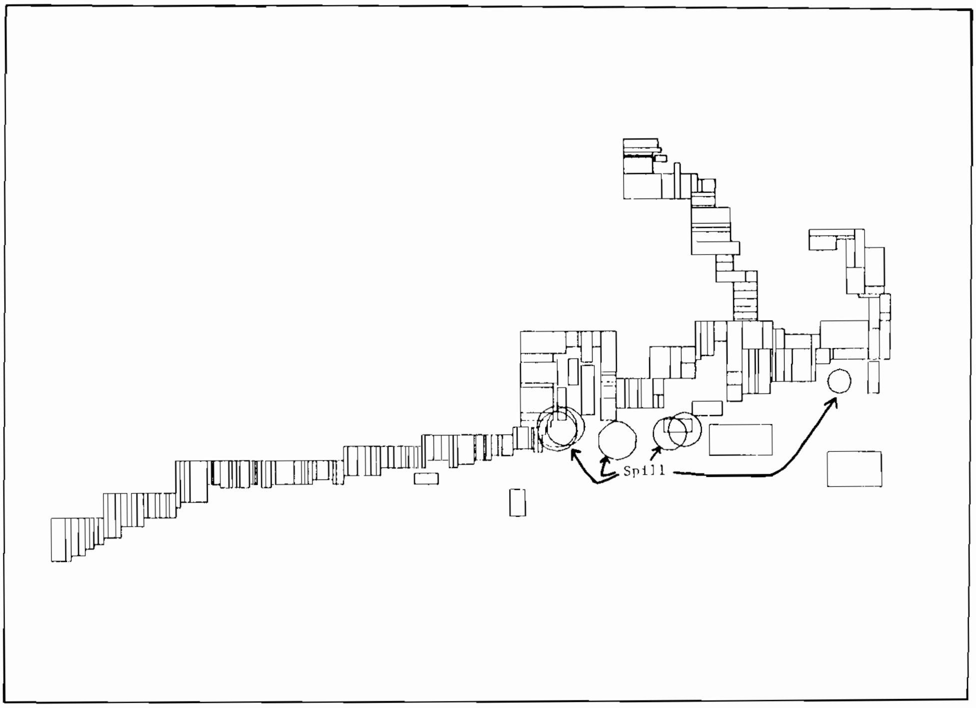

FIGURE 2. *Argo Merchant* Spill/Worst Case Scenario: Impact of spill on Massachusetts and Rhode Island shoreline and island resource grid system.

(1) As expected, damage costs in the worst case scenario were much higher than those in the actual scenario. Actually, in the worst-case scenario, most of the damage was inflicted on natural or economic resources, while in the actual case most of the damage was the value of the lost oil.

(2) In both cases, there was a drastic damage cost reduction when the "do-nothing" response was replaced by the optimal response recommended by the model. In both cases the technique that was given priority was offloading. Other techniques such as mechanical removal were used as backup.

(3) For option (2), the surprisingly lower cleanup cost of the worst-case scenario can be attributed to the significant increase in the operational efficiency of the offloading equipment for Diesel Oil No. 2 in comparison to that for crude oil, due to much lower viscosity.

(4) Offloading seems to be preferred over other techniques for two reasons: First, it has a higher mechanical efficiency (offloading recovers pure oil while mechanical removal can recover only an oil/water mixture from the sea). Second, offloading serves to reduce the overall level of damages by saving the market value of the oil, which would drop to near-zero levels if oil were to be mixed with water.

Still, since weather conditions could render offloading operations difficult, we examined a number of alternative response tactics that did not include offloading. Table 2 summarizes cleanup and damage costs for response options (3), (4) and (5) (worst case scenario only). Both cleanup costs and damage costs of options (3) and (5) are higher than those for option (2). It is also interesting to note that the model recommends to "do-nothing" if only mechanical removal equipment with no storage is available. The absence of adequate storage capacity has frequently proven to be a severe "bottleneck" in cleanup operations and these runs highlight the importance of having *all* components of an equipment package on scene.

Other issues that have been investigated in conjunction with this spill were:

TABLE 2

Summary of Cleanup and Damage Costs for the Worst-Case Scenario and Three Response Options. For Further Explanations See Text.

Response Option	(3)	(4)	(5)
Cleanup Cost	$1.118^M	—	$0.484^M
Damage Cost	$8.34 ^M	$34.8^M	$8.764^M
Total Cost	$8.458^M	$34.8^M	$9.248^M

(1) Sensitivity analysis on notification time (NOTFI) or delays: Although damage and total costs generally increased with longer delays, this was not always the case with cleanup costs. If delays were too long, damages had already occurred and it did not pay to mobilize extensively to prevent them. Such an analysis quantifies the benefits to be derived from an improved monitoring system (one that would reduce the length of the negotiating process among the spiller, the OSC, and the cleanup contractors that precedes the initiation of cleanup).

(2) Sensitivity analysis or damage weight W: One of the most interesting observations *for this spill* was that the response tactic did not change significantly with changes in W. For instance, in the worst-case scenario an underestimation of damage costs by a factor of 20 resulted in the same response tactic. Such "stability" of the problem solution can be very significant whenever there is uncertainty about the value of damages (as there usually is). Thus, it may not always be that crucial to be able to predict the *exact* dollar value of damages in a spill. In many cases, a simple "order of magnitude" estimate of that value may be sufficient to determine the optimal response tactic.

An extensive discussion of those and other runs can be found in Chapter 12 of Ziogas (1982).

5. Discussion

5.1. *On the Damage Function*

This work has assumed a general function $D_n(A_n, d_n, X_n, f_n)$ for the damages that accrue between time stages n and $n + 1$ due to the fact that a portion of oil f_n becomes nonrecoverable during that interval. Any functional form can be used for the damage function, as long as it satisfies the monotonicity assumptions outlined in §3.3 and is timewise separable. One can check that the monotonicity assumptions make sense from a physical point of view. The separability assumption is justified if damages between n and $n + 1$ do not depend on damages that have already occurred up to stage n (that is, if D_n is considered as the *incremental* damage). Special care has been taken within the Damage Assessment Model (Figure 1) so that damages are computed in an unambiguously incremental fashion. It should be stressed again that in addition to A_n, d_n, X_n and f_n, which are explicit variables, the Damage Assessment Model implicitly takes into account variables such as type of oil, spill location, weather conditions, geomorphology and shoreline, environmental and economic resources in the area etc. Damages are broken down into several main categories (value of lost oil, commercial resources, noncommercial resources, beaches, tourism and recreation). The formulation of this paper (which assumes a *single weight* W for the total damage) easily extends to a formulation in which each of the above main categories carries its individual weight. That increases the flexibility of the model by allowing the user to put more emphasis on one or a few damage categories versus others and to investigate the impact of various weighting schemes on the response tactic recommended by the model.

5.2. *Assembling the Equipment into Sets*

A most important aspect of this model is the way in which cleanup equipment is handled within the algorithm. There is no doubt that the way in which the various equipment components are assembled into "equipment sets" is critical for the performance of the system. In addition, since an individual component (such as a pump) may be used in more than one combination with other cleanup components (depending on the particular spill situation), preassigning each component to one and only one set might reduce flexibility in the overall decision-making process. From our discussions with cleanup technology experts, we feel that while such a risk exists in general, its consequences can be greatly offset if the assembly of components into self-contained packages is done intelligently. Past spill history has indicated that it is critical to work with equipment that is self-contained and well-balanced. For instance, combining a 2,000 gallon/hr. pump with a skimmer whose aperture width does not allow a throughput of more than 500 gallons/hr. or which has inadequate storage capacity, is likely to result in an inefficient operation. The identification of "bottlenecks" such as the above is done within the Operational Model (Figure 1). In addition, expert judgement from cleanup contractors and equipment manufacturers should always be an input to the final assembly of equipment into sets, as it has been for all realistic applications of the model thus far. Finally, a measure of last resort (which we do not wholeheartedly recommend) is the following: If it is judged that an individual component could be used in either one of several equipment sets, then the model user has the option of having that component assigned to more than one set, with the understanding that he would have to adjust the problem solution if more than one of these sets are simultaneously selected by the model.

5.3. *Feedback from the Operational Level*

Within the overall model, the most straightforward execution of decisions regarding a specific spill is a one-pass sequence of the form "tactical decisions first, operational decisions second" with no opportunity for feedback. However, since operational decisions (such as what area of shoreline should be protected) may influence the potential or actual level of damages, feedback of *those* damages into the tactical level might be generally warranted. Such a feedback may be in fact necessary if damage computed at the operational level (using an "optimized" spatial allocation of those cleanup resources that have been established at the tactical level) are *significantly* different from those computed at the tactical level (assuming the aggregate response is "uniformly" distributed). Such significant disparities between damages evaluated at those levels are less likely to occur in spills distant from the shoreline, where the main effort is to recover the bulk of the oil before it hits the coast. By contrast, such disparities are more likely in near-shore spills, where the emphasis is on the *protection* of sensitive areas and where spatial allocation issues are more important. Although no feedback mechanisms of such nature have been implemented to date within the overall model, the Tactical Model could be rerun with a new damage weight equal to the original damage weight, times the ratio of the damages computed at the operational level divided by the damages computed at the previous tactical run, *if, in fact, the latter ratio is significantly (i.e. order-of-magnitude) different from* 1.0. The frequent stability of the Tactical Model solutions with respect to changes in the damage weight argues that such iterations between tactical and operational levels would be necessary only rarely (the reader is referred to Ziogas 1982 and Demis 1984 for more on the optimal spatial allocation of cleanup resources in an operational situation).

5.4. *Uses of the Model*

A major use of this model could be for simulation and training purposes. We view such a use as potentially beneficial not only to persons or organizations using the model, but also to the model itself, which could be further improved and refined as a result of such interaction. To make its use more flexible within that context, we have structured the computer program implementing the model into *four* alternative modes of operation:

(a) the "do-nothing" mode, where the model simply tracks down the spill and evaluates damages;

(b) the "manual" mode, where the entire response tactic Y_{in} is entered interactively by the user (nothing is optimized);

(c) the "semiautomatic" mode, where the user enters interactively only the aggregate cleanup capability X_n and the model chooses equipment by solving the corresponding knapsack problem (12); and finally

(d) the "automatic" mode, where the model solves the DP recursion as described in §3. The last mode can function either in a "static" fashion (where no input is updated in time) or in a "dynamic" fashion (where inputs are dynamically updated in time).

Once mastered and refined, the model could be used to assist On Scene Coordinators with their actual responses to spills. A typical use of the model in this mode would involve running the algorithm several times to account for uncertainties in problem inputs, such as weather conditions, oil outflow rate, etc. In such a way the human decision-maker would be able to assess quickly the consequences of various scenarios and response options on a cost/benefit basis.

The *Argo Merchant* application presented in §4 has illustrated yet another potential use of this model: performing "after-the-fact" analyses of oil spill cleanup decisions. Such analyses of both the actual spill scenario and its variants may be used to investigate the merits of response tactics recommended by the model versus those that were actually implemented.

Similar analyses may be used to address "strategic" or "policy" issues. Suppose for instance, that after analyzing a number of serious spills, the model indicated that millions of dollars in damages could have been averted had cleanup operations started six hours earlier than they actually did. Such a result could justify new Federal policies that would shorten the time between a spill and initiation of cleanup operations.

As another example, suppose that a massive spill occurred in adverse weather conditions and the model indicated that a "do-nothing" response would result in \$1 million in damages. Under existing legislation, such a response tactic is both illegal and politically unacceptable. Suppose, however, that the model also indicated that a "do-everything" response, that is, a response that minimized damages irrespective of cleanup costs (weight $W \to +\infty$), would reduce damages by only \$0.1 million and would add \$1 million in cleanup costs. Now there would be a way to assess whether the reduction in damage costs is enough to justify the expenditure for cleanup costs. Given that funds for cleanup operations are by no means unlimited, such exercises of the model could help to shape and justify more efficient spill cleanup legislation, under which pollution combat funds could be spent in more effective ways.

Similar uses of the model could shed more light on other important issues, such as the use of chemical dispersants, the use of private contractors versus Federally owned equipment (which is usually much cheaper), or the question of when to stop cleanup operations (or, "how clean is clean?"). With respect to the last issue, the model clearly shows the diminishing returns associated with prolonging cleanup "until the last drop of oil is removed". Under existing legislation, however, the decision of when to terminate cleanup is not an economic decision.

Overall, it is our opinion that the value of the model developed here lies not so much in its ability to capture the most significant factors of a very complex problem and then solve it, but rather in its ability to answer important "what if" questions so as to provide further insight into the spill cleanup process. The use of the "what if" capability of the model is particularly helpful in problems where reliable input data do not exist. Data reliability is usually a serious problem in all oil spills but is definitely not a reason for not conducting analyses or not running models. Our model can be used to determine which sets of input data are important in the decision-making process (and therefore guide future data collection efforts) and which sets are less important or irrelevant.[1]

[1]Support for the work described in this paper (MIT Oil Spill Model) has been provided by the National Oceanic and Atmospheric Administration, the U.S. Coast Guard, the U.S. Navy, the Commonwealth of Massachusetts, the Doherty Foundation, JBF Scientific Corporation, the Spill Control Association of America, Texaco and Petro-Canada. We would like to thank Warren E. Walker, William A. Wallace and three anonymous reviewers for their comments on two earlier versions of the paper.

References

ANTHONY, R. N., "Planning and Control Systems: A Framework for Analysis," Harvard Business School, Cambridge, Mass., 1965.

BELARDO, S., K. KARWAN, W. WALLACE AND J. WARD, "A Partial Covering Approach to Siting Response Resources for Major Maritime Oil Spills," *Management Sci.*, 30 (1984), 1184–1196.

CHARNES, A., W. COOPER, K. KARWAN AND W. WALLACE, "A Chance-Constrained, Goal Programming Model to Evaluate Response Resources for Marine Pollution Disasters," *J. Environmental Economics and Management*, 6 (1979), 244–234.

CONRAD, J. M., "Oil Spills: The Strategy of Recovery," *Coastal Zone Management J.*, 4 (1978), 409–434.

DEMIS, D. J., "Oil Spill Management: The Damage Assessment Model and the Spatial Allocation of Cleanup Equipment," unpublished M.Sc. Thesis, MIT, May 1984.

DIETZEL, G. F. L., A. W. GLASS AND 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-Blowouts. Its Development and Application to the North Sea," Shell Report EP-47436, 1976.

FRASER, J. P. AND R. A. COCHRAN, "Computer Simulation Aids for Oil Spill Cleanup Decisions," Proc., Offshore Technology Conference, Houston, Tex., 1975.

MILGRAM, J. H., "Being Prepared for Future Argo Merchants," *Technology Rev.*, 79 (1977), 14–27.

POLLACK, A. AND K. STOLZENBACH, "Crisis Science: Investigations in Response to the *Argo Merchant* Oil Spill," MIT Sea Grant Report No. 78-8, 1978.

PSARAFTIS, H. N., "Water Pollution," McGraw-Hill 1982–83 Yearbook of Science and Technology 1982.

———, J. D. NYHART AND D. A. BETTS, "First Experiences with the Massachusetts Institute of Technology Oil Spill Model," Proc., Oil Spill Conference, San Antonio, Tex, 1983.

———, G. G. THARAKAN AND A. CEDER, "Optimal Response to Oil Spills: The Strategic Decision Case," *Oper. Res.*, (to appear).

Seaward International Inc., "Oil Pollution Response Planning Guide for Extreme Weather Conditions," Report to U.S. Coast Guard, 1979.

THARAKAN, G. G., "The Complementary Locations on a Network," unpublished Ph.D. thesis, MIT, June 1982.

TSC, "Deployment Requirements for U.S. Coast Guard Pollution Response Equipment," Transportation Systems Center, Report to U.S. Coast Guard, 1979.

Versar Inc., "Handbook for Oil Spill Protection and Cleanup Priorities," Report to U.S. Environmental Protection Agency, 1981.

WALKER, W. E., J. M. CHAIKEN AND E. J. IGNALL (EDS.), "Fire Department Deployment Analysis: A Public Policy Analysis Case Study," North-Holland, Amsterdam, 1979.

WHITE, I. C. AND J. A. NICHOLS, "The Cost of Oil Spills," Proc., Oil Spill Conference, San Antonio, Tex., 1983.

ZIOGAS, B. O., "Management of Oil Spill Cleanup," unpublished Ph.D. thesis, MIT, February 1982.