

 Open access • Journal Article • DOI:10.1109/21.310528

Operational risk management: a new paradigm for decision making

— [Source link](#) 

Giampiero E. G. Beroggi, William A. Wallace

Institutions: Delft University of Technology, Rensselaer Polytechnic Institute

Published on: 01 Jan 1994 - IEEE Transactions on Systems, Man, and Cybernetics (IEEE)

Topics: Operational risk management, Risk analysis, Risk management, Risk analysis (business) and Risk management information systems

Related papers:

- [Multi-expert operational risk management](#)
- [Operational control of the transportation of hazardous materials: an assessment of alternative decision models](#)
- [A real-time routing model for hazardous materials](#)
- [The Collapse of Sensemaking in Organizations: The Mann Gulch Disaster](#)
- [Identifying, understanding, and analyzing critical infrastructure interdependencies](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/operational-risk-management-a-new-paradigm-for-decision-4qukbiq5us>

Operational Risk Management: A New Paradigm for Decision Making

Giampiero E. G. Beroggi and William A. Wallace

Abstract—The need for more effective ways to manage the risk and safety of technological systems has been widely recognized and accepted by government and industry. Traditionally, risk analysis has been considered as part of the process of planning a technological system and addressed the risk inherent in its day-to-day operations. However, risk must also be considered when responding to episodic events whose uniqueness requires taking actions that are variants of, or different from planned operational procedures. The purpose of this paper is to present a new paradigm for real-time risk analysis that capitalizes upon the advances in computer power, human-machine interfaces, and communication technology. A new process for risk assessment and an appropriate reasoning algorithm for choice has been developed which supports the human operator in analyzing risks and making decisions in real-time during unexpected disruptions in the operations of large-scale systems. The process recognizes that although advances in technology may automate many tasks, humans will always be an integral part of managing large-scale systems. The practical realism of the new approach of operational risk management is illustrated by two examples, hazardous material transportation and emergency management. The first example is discussed within the context of a prototype decision support system for interactive real-time risk management.

Index Terms— Risk assessment, risk management, real-time operational control, ordinal analysis, decision support systems.

I. INTRODUCTION

LARGE-SCALE operations involving technological hazards are typically managed by well formulated, predetermined procedures. These procedures (also called courses of action) are designed to ensure that the operations are as safe and as cost-effective as possible. Unfortunately, large-scale operations are exposed to episodic events that can threaten the successful completion of a planned course of action. Real-time decision making is needed to respond to these unexpected events in order to prevent or mitigate undesirable consequences—in terms of both safety and costs.

Operational decision making has been addressed for various large-scale operational systems, such as air traffic control [1], vehicle navigation [2], mission planning [3], and process control [4], [5]. The methods proposed typically focus on the selection and implementation of planned activities, in some cases, automatically. They do not address the need for revision of these procedures in the light of unanticipated, episodic

events. As noted in a handbook for chemical process safety, “performing high quality hazard evaluations throughout the lifetime of a process cannot guarantee that accidents will not occur” [6]. In addition, emergency procedures themselves are susceptible to unforeseen events, e.g., a “planned for” emergency vehicle is not operational when needed. Also, dealing with the unexpected is still the province of the human decision maker. Therefore, technological aids for supporting decision makers in these typically stressful environments must consider the cognitive limitations of humans in dealing with safety and cost issues in real-time decision making.

Sudden onset events do not necessarily have to affect present operations. A truck carrying hazardous materials could, for example, be scheduled to drive through a region where a snow storm is expected or already in progress. If the dispatcher could “see and think ahead,” he or she could assess the impact of such an event on the planned route and evaluate alternate routes, i.e., avoid the snow storm by changing the planned course of action. Another example is specialized equipment that is to be used in a final stage of an emergency response, but is found to be unavailable. If the on-scene manager could “see and think ahead,” that person could change the planned response to a course of action that does not need that particular equipment.

However, “seeing and thinking ahead,” which is the sensing and reasoning needed for real-time risk analysis and decision making, has only just become possible because the technology to monitor the environment for critical events and communication technology to support real-time reasoning are now available. For example, latest advances in satellite positioning technology and mobile communication make real-time support for hazardous material shipments a feasible approach [7], [8], and advanced monitoring devices, sensors, and forthcoming global mobile communications provide the basis for real-time support for emergency response [9]. These advances are permitting operators to “see” inside a damaged pressure vessel with micro-video, “feel” the conditions inside a burning container with robot devices, and “track” the dispersion of hazardous materials.

This technology provides new capabilities to monitor and control large-scale operations. We would like to capitalize on these advances by a novel approach for supporting real-time risk analysis and decision making for large-scale operations so that pre-planned courses of action can be revised in real-time whenever a critical event occurs. Although advanced technology can automate many tasks, we must consider that humans will always be an integral part of managing large-scale

Manuscript received May 29, 1993; revised December 10, 1993.

G. E. G. Beroggi is with the Department of Systems Engineering, Policy Analysis, and Management, Delft University of Technology, 2600 GA Delft, The Netherlands.

W. A. Wallace is with the Department of Decision Sciences and Engineering Systems, Rensselaer Polytechnic Institute, Troy, NY 12180-3590 USA.

IEEE Log Number 9403038.

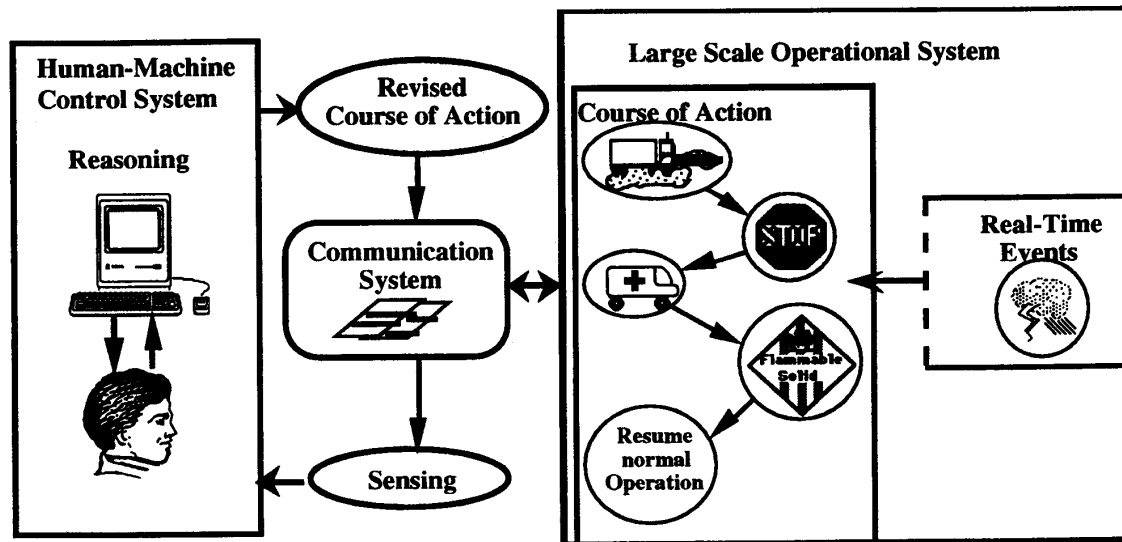


Fig. 1. Environment for managing large-scale operational systems in response to real-time events.

operations and will require decision support in responding to episodic events with the possibility of catastrophic impact.

II. THE REAL-TIME ENVIRONMENT

Risk analysis and decision making in real-time can be described using Anthony's decision taxonomy [10]. *Strategic planning decisions* are the traditional approach in risk management. *Operational risk management* consists of the identification of an unexpected event, an assessment of its consequences and the decision to change the planned course of action.

The real-time environment for managing large-scale systems involving hazardous operations consists of three major components (see Fig. 1): (1) The **large-scale operational system** where operations take place and where events can occur that may affect the planned course of action; (2) the **real-time controller** which senses the operation for accurate and timely performance and the potentially critical events, and reasons about whether the course of action should be revised; and (3) the **communication links** for data transfer between the system and the real-time controller. Due to advanced communication technology, the real-time controller does not have to be located with the operational system. However, we will assume that the loop between sensing and reasoning will always include humans.

Real-time events (RTE's) are occurrences which may force a change in the present course of action. RTE's are unpredictable and occur without warning. Therefore, it is crucial that the human operator senses all RTE's. To do this she or he must continuously monitor the system and its environment. Both the need to reevaluate a course of action and the time available for sensing and reasoning are event driven. A reevaluation is necessary if at least one activity is threatened by the RTE (e.g., the fire truck is useless for response if it is out of order). The time available to make decisions on a new course of action depends upon when the impact of the RTE effects the present course of

action. Since the present course of action is known, this time interval can be estimated. To ensure safe and cost-effective operations of large-scale systems, not only must the appropriate course of action be selected but deadlines must be met.

III. REASONING PRINCIPLES IN LARGE-SCALE OPERATIONAL SYSTEMS

A. The Topological Graph Structure

Reasoning in managing operational large-scale systems consists of assessing the consequences of RTE's on the given attributes (e.g., risks and costs), evaluating alternate courses of action, and deciding to either remain on the present or take a new course of action. Thus, analysis and decision making must be done in real-time. We presume that a human operator cannot both assess effects of RTE's, and generate and evaluate alternate courses of action without support. Information technology can perform computationally intensive tasks in support of human information processing and reasoning. In order to successfully blend human and machine capabilities, a decision structure must be designed. Sage [11] notes that operational decisions are more likely to be made within a set of guidelines than strategic planning decisions. Such a meta-reasoning structure must allow one to assess the impacts and re-evaluate a given course of action whenever a RTE occurs. It also must consider that humans in a real-time decision situation must perform under conditions of uncertainty and time constraints. We suggest using a graph theoretical approach to structure real-time risk analysis and decision making problems.

A course of action consists of a temporal ordered sequence of decisions and concomitant selected activities. Each activity is preceded by a decision, which in turn leads to a new decision on the next activity. Some activities require that other activities be taken previously, others do not. The first decision to be made in operational risk management is to decide whether a




Fig. 2. The topological graph structure.

significant event has occurred and to reevaluate the present planned course of action. The last decision is to decide if "normal" operation can be resumed, i.e., if the hazard is resolved. It is important to note that there are always several different ways to abate the risk for a particular RTE; i.e., there are different courses of actions that can abate the risk for a particular RTE.

Courses of action can be represented on a graph structure. A graph consists of nodes (vertices) and links (edges). Nodes represent decisions and links are activities. An oriented link from decision node D_i to decision node D_j indicates that decision D_i has been made to do activity A_{ij} which leads to decision D_j ; i.e., the link A_{ij} states that there is a relation between the two decisions D_i and D_j . Therefore, every activity leads to a subsequent decision, except for the last decision, i.e., normal operation can be resumed (see Fig. 2).

The graph with oriented links will be called the **topological graph structure** of real-time decision making. A feasible course of action is a sequence of activities on the graph which begins at the first decision node, goes along some intermediate activity nodes (considering the orientation of the links), and ends at one of the end stages (i.e., normal operation). Selecting an activity A_{ij} can also be seen as a walk on the graph from decision node D_i to decision node D_j . A course of action is therefore a walk through the graph from origin (first node) to one of the destinations (last node). Depending on the graph, there might be none, one, or many such walks. The determination of all feasible walks (courses of action) is a combinatorial problem on the topological graph. Therefore, after each RTE, the human operator specifies all feasible walks through the graph (courses of action).

The feasibility of a link (activity) can depend upon external conditions or preceding courses of action. Using a transportation network as an example of a topological graph structure with the nodes being cities and the links connections between two cities, some roads require special permits for shipping hazardous materials. Thus, the feasibility of such roads depends on the cargo that is shipped. In emergency response, the feasibility of links (activities) is in most cases dependent upon previously taken activities. For example, the decision to use chemicals to extinguish a fire can only be made if the appropriate equipment is available.

B. The Preference Graph

Every activity has some expected result. This result, in turn, has a value which is based upon a set of attributes with associated preferences. The set of "best" courses of action

(which is a subset of the set of feasible courses of action) can then be determined by an appropriate algorithmical procedure. The most common attributes in risk management are risks and costs. Risks refer to human life; environmental pollution and destruction; property loss (also expressed in costs); loss of company image; etc. Costs are expenditures for a given activity. We assume that they can be estimated with sufficient accuracy to be considered deterministic.

Risks during operations are typically characterized by high uncertainty and imprecision, especially when evaluated under time constraints and stress. Although operational risks are non-deterministic, no explicit stochastic model can be postulated. Therefore, we consider the possibility of real-time events and their impacts rather than their expectation in a probabilistic sense [12].

We measure the expected results of activities by specifying an ordered set of **preference classes** for the set of attributes. Ordinal measurement means that each attribute must be assigned to one and only one class and that these classes can be ordered. Risk is non-commensurate among classes, i.e., there exists no number n that can measure the risk among classes, but can be measured by an interval scale or cardinally within a class. Therefore, we are assuming that it is only possible to assess whether one activity is more or less dangerous than another one. Expressions like "more" or "less" come from qualitative reasoning [13].

Ordinal relations will not only be used for the assessment of operational risks but also for the tradeoffs between risks and any other attributes, such as costs, benefits, or values. We describe here the procedure of defining the preference classes and their ordinal relations for the case where only the two attributes "cost" and "risk" have been identified.

We assume that the human operator is faced with activities that must be considered as feasible, but have risks so great that there is no monetary equivalent. However, there are also risks that do not justify the expenditures needed to abate or avoid them. We therefore have identified three ordinal classes: High Risks, Costs, and Low Risks.

Two additional classes are also introduced. Activities that do not have any significant risks or costs will be assigned to the first class (ω); activities that under no circumstances are allowed to be engaged in will be assigned to the second class (α). The foregoing classes have the following ordinal relation (the sign "<" means "less preferred"):

$$\alpha < \text{High Risks} < \text{Costs} < \text{Low Risks} < \omega.$$

Since an activity with only low risks and no costs is more preferred than an activity with high risks, the ordinal relation implies a lexicographic preference structure.

Given the topological graph structure, the operator must map the attributes of all activities into the preference classes. Just like with the feasibility of activities (links), the preference of an activity can depend on previous courses of action. For example, using chemical dispersants to respond to an oil spill has a smaller risk (higher safety-preference) if the effectiveness of the dispersants has been previously tested. After all activities have been classified using such a preference structure, the graph is called a **preference graph**.

Whenever a RTE occurs, the human operator must reassess risks and costs of all activities affected by the RTE using the defined preference classes-in real-time. This amounts to reexpressing the preferences of the affected activities considering the RTE. In order to speed up the procedure of reassessment, the human operator does the cognitive part and the computer the algorithmic part. The operator can determine the affected links (activities) by noting them on a graphical display. If the links he or she identifies affect other links in the graph, both the feasibility and preference of these activities must also be reassessed. If, for example, the RTE is "an icy road" then the preference of using a truck to disperse chemicals for fire-fighting will be lower because of the possibility of an accident. Therefore, the preference for using chemicals to fight the fire will also be lower. All the affected activities are presented to the operator one-by-one or grouped into similar operations for reassessment of the preferences. Since the computer does the tedious error-prone search of activities that have to be reassessed, the operator can concentrate fully on the assessment itself. This process of real-time risk analysis is supported by the use of interactive pictorial displays.

After all the affected activities have been reassessed, an algorithm determines the necessary changes in the course of action. The algorithm must resolve two issues: (1) how to "add up" preferences of different activities in order to compute the overall preference of a course of action, and (2) how to determine which of the different feasible courses of action is most preferred. In so doing, it must also consider the limitations of human cognition [14]. According to the definition of preference classes, it is correct to "add up" preferences within a class but not across the classes. In addition, the principle of non-commensurable preference classes must hold throughout the decision process. To answer the second question of how to choose the most preferred course of action (i.e., how to order the feasible courses of action), we will adopt a MinMax strategy which is a well known principle in decision making under uncertainty [15]. The MinMax strategy can be easily incorporated into existing network algorithms.

IV. ILLUSTRATIVE EXAMPLES

A. Transportation of Hazardous Materials

The dispatcher of a truck fleet monitors on a map-like screen the movement of the vehicles [18]. The movement can be seen as a walk through the road network (topological graph structure). Each node is a decision node, since at each node the dispatcher has at least two alternatives to continue the shipment. The dispatcher senses in real-time the environment for weather changes, traffic accidents, traffic congestions, changes in the state of the cargo, etc. The most preferred route under given conditions is depicted in Fig. 3 in dark.

The dispatcher of hazardous material shipments tries to minimize costs without compromising safety. The highest priority is, therefore, to avoid road sections that are too dangerous or too costly (preference-class α). Among the remaining feasible routes, the dispatcher tries not to use roads in the high risk regions. Among the remaining feasible routes,

costs have to be minimized. If any of these routes have the same costs, those with the lowest risk will be selected. The preference structure has therefore the following form:

$$\alpha < \text{Risks}_{\text{high}} < \text{Costs} < \text{Risks}_{\text{slow}} < \omega.$$

All links of the transportation network have been assessed strategically for normal conditions. To do this, the entities along the road network that can have an influence on safety and time, have been identified. Examples are bridges, tunnels, and urban areas. The roads have been divided into six safety-classes, based upon average risk values given in the literature [16]. These classes are highways in urban and rural areas, local roads in urban and rural areas, mountain roads, and road sections in tunnels. The costs have been computed assuming average transportation costs per hour and average travel speed. After all the links have been assessed, optimal routes can be determined for any origin/destination pair on the road network. Under normal conditions, no high-risks were present on the network. Therefore, the optimal routes were determined based solely on transportation costs.

The optimal route for shipment 1 (dark lines in Fig. 3) has the following overall values: cost-preference of 305.69 and 6 low-risk-preferences. The second best route has costs of 307.11 and 4 low risk values. This example shows that in addition to the graphical solution, the overall preference values should also be given. Since the cost- and risk-preferences of the two routes are almost the same, the operator or driver could decide to take the second best route.

While the shipments are on the road, the operator monitors the movement on a computer display, and gathers information about the weather from weather stations and road conditions from the police. Suddenly, the message comes in that a spring thaw caused a large boulder to threaten the underlying road. The operator points on an interactive visual display to the links affected by this event (mouse, touch screen) and the system presents him/her the links one-by-one for reassessment of the risk- and cost-preferences (Fig. 3). For this illustrative scenario, only one link is affected. Typically, more than one link is affected, and the affected links could change the feasibility and preferences of additional links.

After all the affected links have been reassessed, the alternate routes are computed for all the vehicles using an appropriate network search algorithm. The algorithm incorporated into the prototype decision support system in hypermedia [18] is based on the well known shortest path algorithm. The algorithm in this application optimizes preferences using the MinMax principle, instead of the typical minimization of distances. However, not all shipments must have their routes re-evaluated, but only those shipments that planned to drive through the region affected by the event. For this case, the operator changed the risk-preference of the affected link from 1 low-risk to 1 high-risk. The planned route (dark) has now the following values: 1 high-risk-preference, cost-preference of 305.69, and 5 low-risk-preferences. The second best route under normal average conditions is now the optimal route (gray), since it does not have any high-risk-preferences (Fig. 3). As final step, the operator notifies the drivers of the new optimal route in real-time.

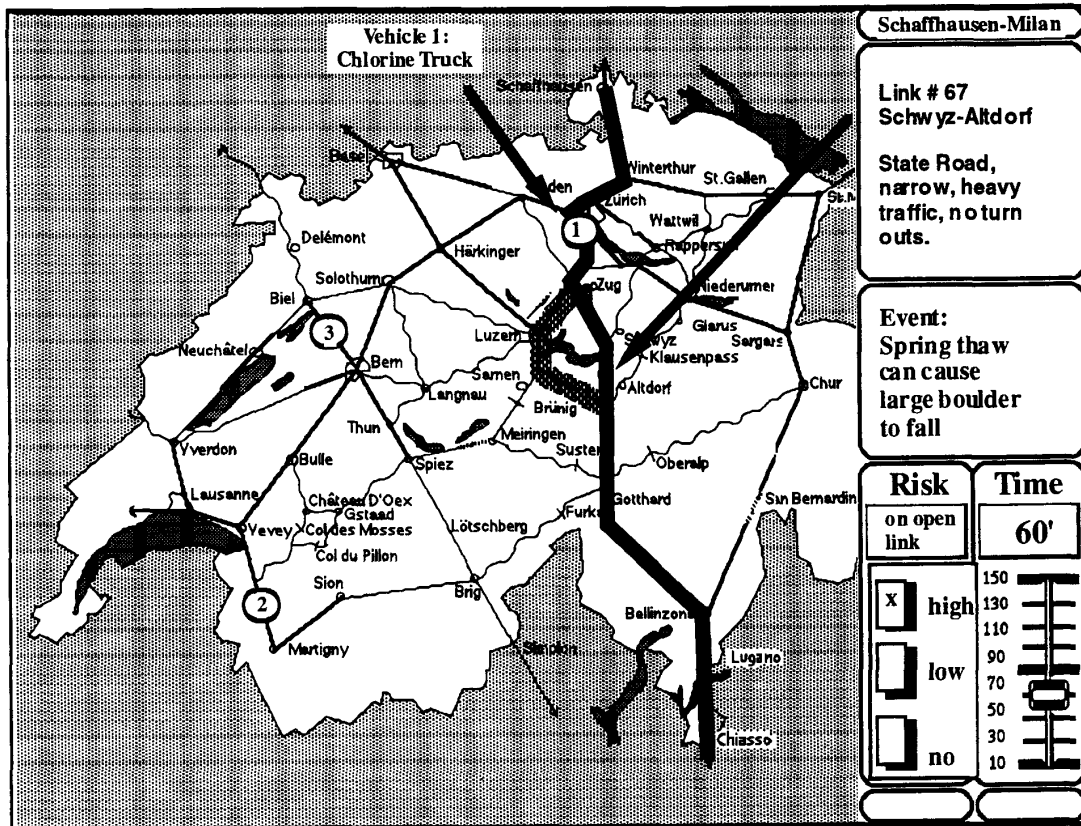


Fig. 3. Shipment number 3 plans to drive through the area affected by the event (dark). By clicking on a link that is affected by the event, a window comes up on the right with the information about the link and the possibility to reassess risk- and cost-preferences. Risk-preferences can be changed by clicking on the appropriate box (link closed, high risk, low risk, no risk), and delays in travel time by moving the slide. After all the links affected by the event have been reassessed, the search algorithm determines for each vehicle the new optimal route (gray) which in turn is transferred to the drivers in real-time.

B. Managing Emergency Response

The second example is based upon the response to the Exxon Valdez oil spill in Alaska, USA [17]. Responding to a maritime casualty where the release of the cargo could have potentially catastrophic impact involves the establishment of a command and control center. The commander, called the on-scene coordinator, and his or her staff gather data, do analyses, make decisions and monitor their implementation. External events, in particular the weather, require that this monitoring and control be done in real-time.

Once an accident occurs, the command and control center is activated. It gathers data on such factors as type of cargo, type of vessel, extent of the accident, and weather, and decides on a course of action—in a very time-constrained situation. The validity of the chosen course of action depends upon the quality of the information available. In addition, the initial decisions act as constraints on subsequent decisions. If, for example, a decision is made not to use chemical dispersants or burning to remove the oil in the water, booms must be allocated to protect vulnerable resources. The response phase is considered complete when there is no threat to human life and the vessel and its cargo are under control.

The attributes chosen for assessing the preference of each action are operational risks (R), operational costs (C), and value of the response activity (V), e.g., the benefit of an activity that mitigates long term environmental damages. The highest priority is given to the safety of personnel on the vessel and in the emergency response team, and the population that may be affected by the release of the cargo. In the case of the Exxon Valdez, the humans at risk were those on the vessel. After the safety of personnel, the emergency response managers must consider environmental damage to the shoreline and the waters, property loss to the shipper, and income loss to fishermen and the tourist industry (now and in the future). The ability of each of the response activities to lesson these impacts is delineated in the class "Values."

An example of the oil spill response activities is shown in Fig. 4. The preference structure was developed using the classes high-risk (h), values (v), costs (c), and low-risk (l), where values are measured by a surrogate for positive impact on the environment, like barrels of oil removed from the water or the shore. It is based upon assessment made in anticipation of the event, which is part of the oil spill response plan. These contingency plans are developed based on an expectation that once the event occurs, response activities can be carried

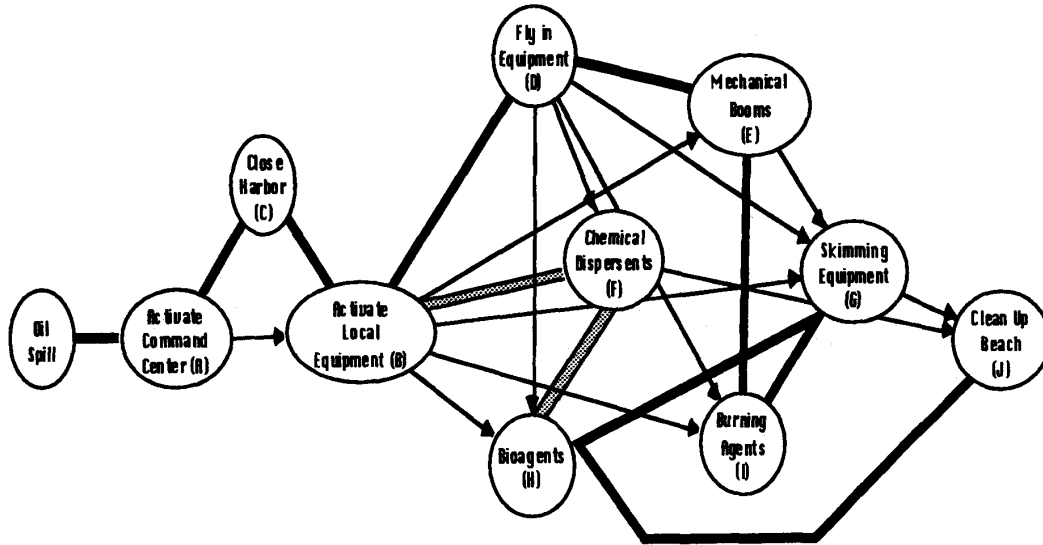


Fig. 4. Topological graph of oil spill response plan. The dark lines represent the optimal course of action before the real-time event. The gray lines show the change in the course of action in response to the real-time event.

out with certainty. It is our contention (and the case in the Exxon Valdez accident) that this expectation is unrealistic, and therefore we need operational risk analysis. The preference structure for the Exxon Valdez response is the following:

$$\alpha < \text{Risks}_{\text{high}} < \text{Values} < \text{Costs} < \text{Risks}_{\text{low}} < \omega.$$

In the previous example, we did not describe in detail the computations needed to arrive at the optimal course of action since the network was too large. For this example, we will illustrate both the assessment of the preferences and the choice process.

The following *a priori* preference relations have been defined, where the letters refer to the circles in Fig. 4, and $\pi(A\mathcal{R}B) = 2c$ means that the preference (π) of the relation (\mathcal{R}) or link between decision A and decision B corresponds to $2c$ (2 units of cleanup costs). We assign preferences for the activities based upon our experience planning for and responding to major oil spills, e.g., Exxon Valdez [17].

$$\{\pi(A\mathcal{R}B) = 2c; \pi(A\mathcal{R}C) = 10c, 1; \pi(C\mathcal{R}B) = c\}$$

$$\begin{aligned} \text{If } A \text{ then: } & \{\pi(B\mathcal{R}D) = c; \pi(B\mathcal{R}E) = 4c; \pi(B\mathcal{R}F) = h, 2v, \\ & 5c; \pi(B\mathcal{R}G) = v, c; \pi(B\mathcal{R}H) = c; \pi(B\mathcal{R}I) = c\} \\ \text{If } A \text{ and } C \text{ then: } & \{\pi(B\mathcal{R}D) = c; \pi(B\mathcal{R}E) = 3c; \pi(B\mathcal{R}F) = \\ & h, 4v, 2c; \pi(B\mathcal{R}G) = v, c; \pi(B\mathcal{R}H) = c; \pi(B\mathcal{R}I) = c\} \end{aligned}$$

$$\{\pi(D\mathcal{R}E) = 8c\}$$

$$\begin{aligned} \text{If } A \text{ then: } & \{\pi(D\mathcal{R}F) = h, 2v, 10c; \pi(D\mathcal{R}G) = v, 5c; \\ & \pi(D\mathcal{R}H) = 2c; \pi(D\mathcal{R}I) = 2c\} \\ \text{If } A \text{ and } C \text{ then: } & \{\pi(D\mathcal{R}F) = h, 8v, 4c; \pi(D\mathcal{R}G) = v, 5c; \\ & \pi(D\mathcal{R}H) = 2c; \pi(D\mathcal{R}I) = 2c\} \end{aligned}$$

$$\text{If } A \text{ and } D \text{ then: } \{\pi(E\mathcal{R}G) = 4v, 5c; \pi(E\mathcal{R}I) = h, v, 2c\}$$

$$\begin{aligned} \text{If } A \text{ and } B \text{ then: } & \{\pi(E\mathcal{R}G) = 2v, c; \pi(E\mathcal{R}I) = h, 3v, c\} \\ \text{If } A \text{ and } C \text{ and } D \text{ then: } & \{\pi(E\mathcal{R}G) = 12v, 5c; \pi(E\mathcal{R}I) = \\ & 4v, c\} \\ \text{If } A \text{ and } C \text{ and } B \text{ then: } & \{\pi(E\mathcal{R}G) = 6v, c; \pi(E\mathcal{R}I) = v, c\} \end{aligned}$$

$$\begin{aligned} \text{If } A \text{ and } B \text{ and } E \text{ then: } & \{\pi(I\mathcal{R}G) = h, 2v, c\} \\ \text{If } A \text{ and } C \text{ and } B \text{ and } E \text{ then: } & \{\pi(I\mathcal{R}G) = 4v, c\} \\ \text{If } A \text{ and } B \text{ then: } & \{\pi(I\mathcal{R}G) = h, v, c\} \\ \text{If } A \text{ and } C \text{ and } B \text{ then: } & \{\pi(I\mathcal{R}G) = v, c\} \\ \text{If } A \text{ and } D \text{ and } E \text{ then: } & \{\pi(I\mathcal{R}G) = h, 2v, c\} \\ \text{If } A \text{ and } C \text{ and } D \text{ and } E \text{ then: } & \{\pi(I\mathcal{R}G) = 8v, c\} \\ \text{If } A \text{ and } B \text{ and } D \text{ then: } & \{\pi(I\mathcal{R}G) = h, v, c\} \\ \text{If } A \text{ and } C \text{ and } D \text{ then: } & \{\pi(I\mathcal{R}G) = 2v, c\} \end{aligned}$$

$$\{\pi(G\mathcal{R}H) = v\}$$

$$\begin{aligned} \text{If } I \text{ then: } & \{\pi(G\mathcal{R}J) = v, 10c, 1\} \\ \text{If } \neg I \text{ then: } & \{\pi(G\mathcal{R}J) = v, 2c, 1\} \end{aligned}$$

$$\{\pi(F\mathcal{R}J) = c, 1\}$$

$$\begin{aligned} \{\pi(F\mathcal{R}H) = l\} \\ \text{If } D \text{ and } G \text{ then: } & \{\pi(H\mathcal{R}J) = 8v, c\} \\ \text{If } D \text{ and } \neg G \text{ then: } & \{\pi(H\mathcal{R}J) = 4v, c\} \\ \text{If } \neg D \text{ and } G \text{ then: } & \{\pi(H\mathcal{R}J) = v, c\} \\ \text{If } \neg D \text{ and } \neg G \text{ then: } & \{\pi(H\mathcal{R}J) = v, c\} \end{aligned}$$

$$\text{If } F \text{ and } \neg D \text{ then: } \{\pi(H\mathcal{R}J) = v, c\}$$

The best response for this particular spill situation, i.e. the optimal path through the graph, is to close the harbor, fly in additional equipment, use booms, use burning agents (which requires the use of skimmers), and add bioagents, as shown in Fig. 4. The solution was determined by first enumerating the

- York, p. 9, 1992
- [7] E. K. Morlok and S. F. Hallowell, "Reported benefits of advanced vehicle tracking and communications systems," *Mobile Satellite Project WP 89-8-1*, Department of Systems Engineering, University of Pennsylvania, 1989.
- [8] G. E. G. Beroggi and W. A. Wallace, "Closing the gap-transit control for hazardous material flow," *Journal of Hazardous Materials*, 27, 1991 (Special Issue), pp. 61-75.
- [9] W. A. Wallace, "On managing disasters: the use of decision aid technology," in: A. Kirby (Ed.), *Nothing to Fear: An Examination of Risks and Hazards in American Society*, Tempe, AZ: University of Arizona Press, 1990.
- [10] R. N. Anthony, *Planning and Control Systems: A Framework for Analysis*, Cambridge, MA: Harvard University Press, 1965.
- [11] A. P. Sage, "Information technology for crisis management," *Large Scale Systems*, 11, pp. 193-205, 1986.
- [12] J. S. Ostroff, *Temporal Logic for Real-Time Systems*, Somerset, UK: Research Studies Press Ltd., 1989.
- [13] J. Kalagnanam, H. A. Simon and Y. Iwasaki, "The mathematical bases for qualitative reasoning," *IEEE Expert* (forthcoming).
- [14] D. Kahneman and A. Tversky, "Variants of uncertainty," *Cognition*, No. 11, pp. 143-157, 1982.
- [15] M. D. Resnik, *Choices*, Minneapolis, MI: University of Minnesota Press, 1987.
- [16] F. F. Saccomanno and A. Y.-W. Cha, "Economic evaluation of routing strategies for hazardous road shipments," *Transportation Research Record*, vol. 1020, pp. 12-18, 1985.
- [17] J. Harrauld, H. Marcus and W. A. Wallace, "The Exxon Valdez: An assessment of crisis prevention and management systems," *Interfaces* vol. 25, no. 5, pp. 14-30, 1990.
- [18] G. E. G. Beroggi and W. A. Wallace, "A prototype decision support system in hypermedia for operational control of hazardous material shipments," *Decision Support Systems* (forthcoming).
- [19] G. E. B. Beroggi and W. A. Wallace, "Operational control of the transportation of hazardous materials: an assessment of four decision models," *Management Science*, (forthcoming).
- [20] G. E. G. Beroggi and W. A. Wallace, "A decision logic for operational risk management," in Carley K. and Prietula M. (eds.), Hillsdale, NJ: Lawrence Elbaum Associate, 1994.



Giampiero E. G. Beroggi is an Assistant Professor of Policy Analysis at the School of Systems Engineering, Policy Analysis, and Management, Delft University of Technology, The Netherlands. He received his Ph.D. and M.S. (OR&Stat) from Rensselaer Polytechnic Institute, Troy, NY, and his Dipl. Ing. from ETH-Zurich, Switzerland. His research interests include operational control of large scale systems, risk and environmental management, decision support systems, and multicriteria models.



William A. Wallace is Professor in Decision Sciences and Engineering Systems at Rensselaer Polytechnic Institute, Troy, NY. He has over 20 years experience in research and development in management science and decision support systems, with particular emphasis on crisis management. His recent work includes a report for the National Science Foundation on the crisis prevention and management systems used in the organized response to the Exxon Valdez oil spill.