Crisis Response Planning: A Task Analysis

Wayne Iba and **Melinda T. Gervasio**
Institute for the Study of Learning and Expertise
2164 Staunton Court
Palo Alto, California 94306
{Iba,Gervasio}@isle.org

Abstract

The ability to respond effectively and efficiently to crisis is an important objective in many problem-solving domains. In this paper, we identify three central elements of crisis: threat, urgency, and uncertainty. We define the space of competent problem solvers for crisis response through a task analysis (Newell and Simon, 1972). Specifically, we bound the space of problem solvers by determining the features of problem spaces and search control strategies necessary to address the central elements of crisis. We also outline an evaluation method for problem solvers in crisis domains. The task analysis provides a framework for improving crisis response strategies as well as developing intelligent assistants for crisis response.

Introduction

Consider the following crisis scenario set on the small island nation of Pacifica, a U.S. territory.[1] A volcano has just erupted on Pacifica and further eruptions and earthquakes endanger the lives of the inhabitants of Barnacle, Calypso, and Delta. The U.S. military decides to evacuate the inhabitants off the island. Calypso and Delta each have a seaport and an airport, allowing evacuation by sea or by air to either Avalon or Honolulu.

The planning task associated with this scenario involves many decisions. In addition to moving the inhabitants off the island, the evacuation operation will also involve bringing personnel and equipment to the island—e.g., military troops, medical aides, trucks, tents, etc. Aside from the different ports and roads, there are different types of sea, air, and ground transport involving different tradeoffs. For example, the B747 is a large plane capable of carrying 408 passengers and traveling at 450 mph, while the C-130 is a small plane capable of carrying only 8 passengers and travelling at 280 mph. However, the C-130 is small enough to land at any airport, while the B747 requires long runways, available only at the Delta and Honolulu airports. The different types of transport also involve different times for loading, transit, and unloading. Because multiple trips must be made to evacuate the island, there are choices in the ordering and methods of evacuation. There are also constraints on the available resources. For example, concurrent military exercises in the Atlantic may make particular units unavailable for the evacuation operation. Underestimated travel times and vehicle breakdowns may reduce or delay resources on which a plan depended.

Crisis response is an issue of great concern to planners in various fields where significant amounts of highly valued assets can be at risk—e.g., relief organizations, business, the military. Cognitive science brings a unique perspective to the treatment of crisis response that is of potential value to organizational planners. This paper presents a task analysis of crisis response with the objective of defining the space of competent problem solvers. This task analysis can be used to understand how human planners perform in crisis situations and their cognitive limitations with respect to the demands of the task. The analysis can also be used as a "requirements specification" for the design of a crisis response assistant that is tailored to improve the level of human performance. In the following sections we present the crisis response task and then the task analysis proper. We briefly describe how problem solvers (both human and artificial) can be evaluated in the crisis response domain and close with a discussion of related and future work.

The Crisis Response Task

The first step in a task analysis is to understand the nature of the task involved—in our case, the crisis response task. At an abstract level, crisis response is simply a planning problem. However, crises have distinguishing features that constrain the problem in interesting ways; these consequently constrain the applicable problem solvers. We identify these distinguishing features of crisis and use them in presenting a formal definition of the crisis response task.

Elements of Crisis

The crisis literature, predominantly in international relations and business management, provides various definitions of crisis (e.g., (Rosenthal et al., 1989; Milburn, 1972)). However, three themes—threat, urgency, and

[1] The setting is the Precis environment (Reece et al., 1993), slightly modified for simplicity.

uncertainty—prevail across these definitions.[2]

The first element of crisis is *threat*. In a crisis, something of great value to the agent is at risk—a house, a family member, money, power, life, etc. The world is progressing in such a way that without intervention, the agent will lose this thing of great value. The agent is thus compelled to act. In the evacuation scenario, the inhabitants' lives are at risk. If they remain in Pacifica, they are likely to perish from direct or indirect effects of the volcanic eruption. However, the threat to their safety may be eliminated by evacuating them off the island.

An interesting consequence of threat is that the primary goal in crisis response planning is a goal of prevention rather than a traditional planning goal of achievement. This potentially adds complexity to the problem in that equivalent goals of achievement may have to be determined for planning purposes. In the evacuation scenario, the goal of preventing death was made equivalent to the goal of relocating the inhabitants of Pacifica to Honolulu and Avalon. However, it might instead have been equated to the goal of containing the volcanic debris or the goal of housing all the inhabitants in fortresses.

The second element of crisis is *urgency*. Urgency refers to the limited time available to an agent in which to respond to a crisis. Specifically, the agent must act before it loses the threatened value. Because time is limited, the agent may be unable to consider all possible options. Within this time, the agent must also not only develop a course of action but implement it as well. This limits the useful options available to the agent. In the evacuation scenario, the island must be completely evacuated before the eruption and its after-effects destroy the roads and ports. Evacuating Barnacle is most urgent, as it has no ports and the lava flow threatens to overrun its only access roads. For this purpose, the only useful trucks are those known to be available before Barnacle becomes isolated.

The final element of crisis is *uncertainty*. Uncertainty takes on various forms and pervades the whole crisis response task. For example, faulty or limited sensors provide noisy information that leads to uncertainty about the current state of the world. Exogenous events due to other agents capable of manipulating the environment result in changes in the world state without the participation, and possibly without the knowledge, of the agent. Actions with stochastic outcomes lead to an agent's inability to precisely predict the effects of action execution. In the evacuation scenario, it is known that Barnacle will eventually be unreachable, but it is not known precisely how long it will take the lava flow to reach the roads. Subsequent eruptions, rain, and various other conditions may factor into the equation. Similarly, approximate times for vehicle loading, transit, and unloading are known, but the actual operating time may vary due to adverse weather conditions, vehicle breakdowns, inhabitants refusing to leave their homes, etc.

[2]Robert Steinrauf, personal communication.

Task Definition

A crisis is necessarily situated within an agent—a value of the agent that is at risk. We now define the scope of the task faced by such an agent in crisis.[3] Given a crisis situation, characterized by the elements of threat, urgency, and uncertainty, the response task involves developing a course of action to eliminate the threat, within the time available, given uncertain information about the changing environment. We define the crisis response task as follows.

> **Crisis Response Task.** Given a tuple $\langle E, A, I, V \rangle$ denoting a crisis situation, determine a plan P that will resolve the crisis.

A *crisis situation* $\langle E, A, I, V \rangle$ defines particular relations between the environment (E), the actions available to the agent (A), the information the agent has about the world (I), and the values of the agent (V). These relations embody the central elements of crisis. The environment E threatens some value $v \in V$ of the agent. Furthermore, E is progressing such that the agent has limited time in which to resolve the threat, given its available actions A. The agent's information I about the world and its knowledge about its actions A is at best uncertain. Given such a situation, the agent must come up with its response—a plan. We use the term *plan* to refer to a knowledge structure that guides the decisions on action execution.

A Task Analysis of Crisis Response

In the previous section we discussed the nature of crisis and defined a crisis situation. In the language of Newell & Simon (1972) we explored the characteristics of the *task environment* and discussed its ramifications for a problem solver. We now consider these ramifications in greater detail by exploring the space of problem solvers for crisis response planning.

A task analysis bounds the space of possible problem solvers for some problem or task. Based on invariants in the task environment, we can bound the set of imaginable problem solvers that might solve a particular task. For our purposes, we will consider only those problem solvers that are plausibly competent for the crisis response task. Thus, we will focus on those problem solvers with means of addressing the threat, urgency, and uncertainty that characterize a crisis.

A problem solver may be defined as the combination of a state space, search operators to explore this space,

[3]In general, crisis response planning may involve multiple agents, thereby raising communication and coordination issues. However, the individual agent still faces a planning problem, albeit additionally constrained. To better focus the discussion in this paper, we analyze only the individual agent.

and search control strategies to guide the exploration (Newell and Simon, 1972). The state space consists of the set of all knowledge states where a state consists of I,A, and V (the information the agent has about the environment, its actions, and its values) and is possibly augmented with internal memory structures (e.g., history, concepts, plans, etc.). The search operators determine the connectivity of the state space by defining the allowable transformations on these knowledge states. These transformations either modify memory or "simulate" action execution. Through these transformations, the problem solver can reason about the task and formulate a plan or sequence of actions that will convert the current state of the environment into a desired state that resolves the crisis. Finally, search control strategies guide a problem solver in its navigation and exploration of the problem space.

Knowledge Representation

The three themes guiding this discussion are the threat within the crisis, the urgency to act so as to avert the threat, and the uncertainty that complicates the crisis response task. These unique characteristics of crisis mandate features and capabilities of possible problem solvers that address the passage of time and the uncertainty of both action outcomes and environmental conditions. In the following sections, we describe the components of a problem solver in light of the constraints imposed by these themes of crisis.

Goal state The problem solver must represent its goals or values (V) within the problem state in a manner that allows the solver to determine when the goal has been achieved. Traditionally, the goal of a planner or problem solver has been framed as a set of desired conditions to be achieved in the environment. Actions are most easily conceptualized in terms of the conditions that they establish or the changes they cause in the environment—not in terms of the conditions that they do not establish. In crisis, however, the primary goal is one of prevention—specifically the prevention of losing something of great value—rather than of achievement.

Fikes, Hart & Nilsson (1972) propose using negated conditions to solve *negative goals*. However, achieving the negation of a condition to be prevented (assuming such negated conditions are represented in some operator's effects to begin with) is in general an insufficient approach to prevention. A goal of prevention is usually associated with a temporal scope; that is, the goal is to prevent a condition from becoming true throughout some duration of time. For this reason, goals of prevention have been addressed primarily within frameworks that provide some means for temporal reasoning.

There are two general approaches to representing and reasoning about goals of prevention. Both distinguish goals having temporal scope from those without (i.e., goals of achievement). The first involves a direct treatment of goals of prevention—through inference rules and operators that directly reference goals of prevention (McDermott, 1982) or through prevention techniques incorporated into the problem-solving mechanism (Hogge, 1988). The second involves transforming goals of prevention into equivalent goals of maintenance (e.g., the goal of preventing the condition of getting wet is the same as the goal of maintaining dryness) and using existing techniques for handling goals of maintenance (Kaelbling, 1988; Schoppers, 1987).

Information from the environment One of the components of a crisis situation is the information about the environment (I) that is available to the problem solver. For the crisis response task, we assume that I includes attributes about the world such as time of day and locations and quantities of relevant resources. This information may also include trends occurring in the environment but the problem solver may be required to infer such trends. Two of the themes of crisis, uncertainty and urgency, impact what information is necessary to a solution and constrain how that information is represented.

Because we are interested in the crisis response task as humans face it in the real world and in tools that facilitate their responses, we must accept that the problem solver's information about the world is uncertain. Sensory information is subject to incompleteness and noise. More importantly, the environment changes over time because exogenous events or other agents can generate surprises. Therefore, a crisis response problem solver must be able to construct a response in spite of uncertain information and multiple possible futures.

One approach to moderating uncertainty from incomplete information is to augment the representation of the knowledge state by actively collecting additional information not initially available in I. However, there are usually costs in active sensing and the search control strategy (discussed later) must use this approach selectively (Chrisman and Simmons, 1991; Draper et al., 1994).

Uncertainty arising from multiple possible futures can be addressed by explicitly representing uncertainty within the problem solver's knowledge state. States may be associated with probabilistic values and the transitions to future states may be represented as probabilistic transitions (Drummond and Bresina, 1990; Hanks, 1990; Dean and Kanazawa, 1988).

Aside from handling uncertainty, a useful or competent crisis response system must also address urgency. Although we have touched on the environment changing over time in the context of uncertain future states, time becomes a critical issue in the context of urgency in that something bad will happen *soon*. During the problem solving process, the problem solver must be able to reason about the current time, the imminence of the threat, and the expected time to execute the actions that will circumvent the crisis. Urgency seems to require that a problem state be time dependent in some manner; that is, the problem state must have the representational expressiveness to describe a changing world.

One approach to dealing with the changing environment is to include a time value in the problem state. This can effectively be a time-stamp on the problem state or a time window during which an action executes. A more comprehensive approach to dealing with time is to use a temporal reasoning framework (Allen, 1984; McDermott, 1982; Kanazawa and Dean, 1989) to allow the problem solver to represent and reason about action durations and orderings.

Actions As part of the crisis situation input, the actions (A) that are available at a given time must be representable within the problem solver's knowledge structure. We assume that an action is executed in the environment, that it has some preconditions that must be satisfied, that it takes some amount of time to execute, that it has some cost and may consume resources (other than time), and that it has the effect of changing certain specified conditions in the environment. Just as a problem solver cannot perfectly represent its environment, it cannot perfectly predict future states of the environment because of uncertainty. Actions effectively behave nondeterministically and must therefore be represented and used accordingly.

One approach to addressing uncertainty is to augment the typical actions (for manipulating the environment) with sensor actions (for gathering information about the environment). These actions also have an associated cost and their outcomes may also be uncertain. However, when utilized appropriately by the problem solver, they can improve the reliability and quality of a solution.

Another approach is to represent actions as having multiple possible outcomes (Warren, 1976; Peot and Smith, 1992), possibly associated with probabilities (Drummond and Bresina, 1990; Draper et al., 1994). This extension is analogous to (and a special case of) the multiple possible futures issue discussed earlier. The execution of a particular action may have some desired or intended effect with some probability, while various other outcomes are also possible with varying probabilities.

Another approach to uncertainty is to apply learning methods to acquire more accurate models of the environment and of the available actions (Gil, 1994; Sutton, 1990). Obviously this approach is only useful if the initial estimates are poor or missing altogether and the cost of learning is acceptable. Furthermore, even after learning the "correct" probabilistic model the problem solver must still deal with the remaining uncertainties.

Problem Space Operators

At this point we must distinguish between actions that the problem solver may execute in the environment and *operators* that manipulate the problem space. The directed links between states in the problem space correspond to operators that the problem solver can apply to reason about the problem and construct an appropriate response. Because these operators determine the connectivity of the problem space, their nature is strongly dependent on the specific knowledge representation. However, without loss of generality, we can assume that a problem state represents the "plan in progress"—i.e., the basic (required) operators correspond to augmenting or repairing some existing (possibly empty) plan. For example, an operator could complete a partial plan by concatenating additional plan fragments. In the case of a reactive problem solver that learns, where a plan consists condition-response rules, operators could either modify the applicability conditions of the reactive rules or add new rules altogether.

In our discussion of the knowledge representation, we introduced probabilistic action outcomes as an approach to dealing with uncertainty. If we assume such a representation, corresponding operators would modify the probabilities on possible outcomes (based on experience) or perhaps add (or delete) contingency plans for possible (or impossible) outcomes. Together with the action representation, these operators enable a problem solver to deal with the issue of uncertainty.

The issue of urgency may be addressed more directly by operators. Because operators define the connectivity of the problem space, higher-order operators can allow a problem solver to explore the space more quickly. For example, in representations where operators correspond to actions, one such higher-order operator is a macro-operator that combines the effects of several actions (Fikes et al., 1972). Another is an abstract operator that represents a class of related actions (Sacerdoti, 1974; Tate, 1977). The indexing and adaptation operators of a case-based planning system (Hammond, 1986) are other examples of higher-order operators that allow larger leaps in the problem space. By reducing the time required to construct an appropriate response, these higher-order operators effectively mitigate the level of urgency faced by the problem solver. However, such operators must be used effectively or they can end up making search more expensive. Regardless of the knowledge representation and operators adopted, the problem solver must still *effectively* search the resulting problem space.

Search Control

The problem solver uses a search control mechanism to select among alternative operators in the process of selecting an output response. The search control mechanism must manage its time by deciding when to stop planning and generate the output. The output of a crisis response planner may take on various forms—a complete plan, a single action, a query for additional information, even a directive to do nothing at all. Possible stopping criteria include running out of time, finding a goal state, needing crucial information, an external interrupt demanding an output, or detecting a significant change in the world state that renders the current plan obsolete.

Search control can be used to address the urgency of a crisis in two ways. First, an efficient search control mechanism can mitigate the urgency of a crisis by guiding the search to a goal state quickly, with little wasted search. Learning effective search control rules (Minton, 1988) is a possible approach to efficient search control strategies. Second, search control can incorporate mechanisms for directly controlling deliberation time subject to externally imposed constraints. For example, anytime planning algorithms allow an agent to come up with an optimal plan for any amount of deliberation time (Boddy and Dean, 1989; Drummond and Bresina, 1990).

Because the goal of the search is to find a state that resolves the threat—ideally, a goal that corresponds to an optimal plan—the issue of threat is also implicitly addressed by search control. Decision-theoretic methods may be incorporated into search control for the purpose of finding such goal states (Feldman and Sproull, 1977).

Search control can also be used to address uncertainty. Specifically, feedback from the environment during execution can help mitigate the problem of incomplete or incorrect initial information. For example, conditional plans that incorporate tests (information gathering actions) at appropriate points throughout the plan and provide alternative branches of the plan according to the expected or possible outcomes of the tests (Peot and Smith, 1992; Draper et al., 1994). This approach has the advantage of anticipating eventualities but consequently can pay a significant cost in planning. Another approach is to interleave planning and execution to construct plan fragments that are likely to complete successfully and then to reassess the environment and plan from the new current state. This approach avoids the up-front computational cost of conditional planning but can be subject to costly dead ends because of the approach's limited look ahead.

In summary, we have identified bounds on possible designs that have a reasonable chance of being competent crisis response problem solvers. We identified research efforts that address various issues we have raised; none of them address all of the issues. Ultimately, new problem solvers motivated by further research must be tested empirically against each other and against human crisis response planners.

Evaluation

There are three primary dependent measures of interest: world-state outcome, timeliness, and cost. The first is the state of the environment that arises from executing the sequence of actions recommended by the problem solver. This is mostly concerned with averting the original threat but could include qualitative factors that value one goal state above another. A second measure is the timeliness of the action recommendations and the successful environmental outcome. We consider a problem solver that resolved a crisis with time to spare to be superior to one that just made it under the wire. Finally, actions, conditional tests, and planning to act all incur hidden or direct costs. By accumulating these costs for a single crisis response task, we can compare problem solvers and prefer those with lower costs.

Naturally, there are many variables that we could manipulate as independent measures. However, in keeping with our focus on crisis, we would expect to vary the three central components of crisis: threat, urgency, and uncertainty. We predict that we will find interactions between these dimensions. That is, a moderate level of uncertainty and urgency may lead to poorer performance than a low level of urgency and a high level of uncertainty (and vice versa). Uncertainty is easily manipulated by controlling the environmental consequences of actions. Urgency can be increased by simply giving the problem solver less time in which to find a solution. Threat may be more problematic to vary but one approach would be to vary the number of values that are attacked (threats) for a given suite of goals. We predict that performance will deteriorate as the severity of the threats increases.

Finally, we cannot stop after evaluating the artificial problem solvers. We must compare their performance (both strengths and weaknesses) to that of human crisis managers. For the most part, this can be accomplished by placing human experts in identical situations as concocted for the artificial problem solvers and manipulating the independent variables accordingly.

Related and Future Work

There have been a few AI planning systems that have been constructed specifically for crisis response planning. Using the framework presented in this paper, we now briefly discuss two that have dealt specifically with noncombatant evacuation operations, as in the Pacifica scenario. O-Plan (Tate et al., 1994) is a hierarchical planning system that has been applied to many large-scale planning tasks. In O-Plan, the issue of threat is addressed through the use of standard operating procedures (SOPs) to plan for user-defined tasks. Urgency is addressed through hierarchical planning as well as scheduling capabilities for reasoning about resource constraints, including time constraints. An associated reactive execution component mitigates the effect of uncertainty by allowing the system to adapt executing plans to the runtime situation. In (Bienkowski and desJardins, 1994), a generative planning system is augmented with components for temporal reasoning, case-based reasoning, and scheduling. The temporal reasoning component addresses the temporal aspect of threat, while the case-based reasoning and scheduling components help to mitigate urgency. Because its primary focus was to speed up crisis response, this work did not address execution and the associated uncertainty.

In identifying the necessary features of problem spaces and search control strategies for crisis response problem solvers, we have focused on requirements specification. With our goal of developing intelligent crisis response assistants, the next step is to focus on the interactions between these requirements and the cognitive capacities of human crisis response planners. Given the importance of crisis response to many real-world domains, we believe more comprehensive analyses of the crisis response task will be invaluable. We hope others will build upon the groundwork we have laid here and explore issues such as coordination and cooperation among multiple agents to complete the task analysis.

Acknowledgements

The Office of Naval Research supported this research through grant number N000014-96-1-1221. Our conceptualization of crisis benefitted from a literature survey by Robert Steinrauf and discussions with Professor Michael Fehling's research group in the EES/OR Department at Stanford University. We thank Pat Langley for useful discussions on the task analysis.

References

Allen, J. F. (1984). Towards a general theory of action and time. *Artificial Intelligence*, 23.

Bienkowski, M. A. and desJardins, M. E. (1994). Planning-based integrated decision support systems. In *Proceedings of the Second International Conference on Artificial Intelligence Planning Systems*, pages 196–201.

Boddy, M. and Dean, T. (1989). Solving time-dependent planning problems. In *Proceedings of IJCAI89*, pages 979–984.

Chrisman, L. and Simmons, R. (1991). Sensible planning. In *Proceedings of AAAI91*, pages 756–761.

Dean, T. and Kanazawa, K. (1988). Probabilistic temporal reasoning. In *Proceedings of AAAI88*, pages 524–528.

Draper, D., Hanks, S., and Weld, D. (1994). Probabilistic planning with information gathering and contingent execution. In *Proceedings of the Second International Conference on Artificial Intelligence Planning Systems*, pages 31–36.

Drummond, M. and Bresina, J. (1990). Anytime synthetic projection: Maximizing the probability of goal satisfaction. In *Proceedings of AAAI90*, pages 138–144.

Feldman, J. and Sproull, R. (1977). Decision theory and AI II: The hungry monkey. *Cognitive Science*, 1:158–192.

Fikes, R. E., Hart, P. E., and Nilsson, N. J. (1972). Learning and executing generalized robot plans. *Artificial Intelligence*, 3(4):251–288.

Gil, Y. (1994). Efficient domain-independent experimentation. In *Proceedings of the Tenth International Conference on Machine Learning*, pages 128–134.

Hammond, K. (1986). Chef: A model of case-based planning. In *Proceedings of AAAI86*, pages 267–271.

Hanks, S. (1990). Practical temporal projections. In *Proceedings of AAAI90*, pages 158–163.

Hogge, J. (1988). Prevention techniques for a temporal planner. In *Proceedings of AAAI88*, pages 43–48.

Kaelbling, L. P. (1988). Goals as parallel program specifications. In *Proceedings of AAAI88*, pages 60–65.

Kanazawa, K. and Dean, T. (1989). A model for projection and action. In *Proceedings of IJCAI89*, pages 985–990.

McDermott, D. (1982). A temporal logic for reasoning about processes and plans. *Cognitive Science*, 6:101–155.

Milburn, T. W. (1972). The management of crises. In Hermann, C. F., editor, *International Crises: Insights from Behavioral Research*. Free Press, New York.

Minton, S. (1988). *Learning Search Control Knowledge: An Explanation-Based Approach*. Kluwer Academic Publishers, Norwell, Massachusetts.

Newell, A. and Simon, H. A. (1972). *Human Problem Solving*. Prentice-Hall, Englewood Cliffs, New Jersey.

Peot, M. A. and Smith, D. E. (1992). Conditional nonlinear planning. In *Proceedings of the First International Conference on Artificial Intelligence Planning Systems*, pages 189–197.

Reece, G. A., Tate, A., Brown, D. I., Hoffman, M., and Burnard, R. E. (1993). The PRECiS environment. In *Proceedings of the ARPA-RL Planning Initiative Workshop*.

Rosenthal, U., Charles, M., t'Hart, P., Kouzmin, A., and Jarman, A. (1989). Case studies to theory and recommendations: A concluding analysis. In Rosenthal, U., Charles, M., and t'Hart, P., editors, *Coping with Crises*.

Sacerdoti, E. (1974). Planning in a hierarchy of abstraction spaces. *Artificial Intelligence*, 5:115–135.

Schoppers, M. J. (1987). Reactive reasoning and planning. In *Proceedings of AAAI87*.

Sutton, R. S. (1990). Integrated architectures for learning, planning, and reacting based on approximating dynamic programming. In *Proceedings of the Seventh International Conference on Machine Learning*, pages 216–224.

Tate, A. (1977). Generating project networks. In *Proceedings of IJCAI77*.

Tate, A., Drabble, B., and Kirby, R. B. (1994). O-Plan2: An open architecture for command, planning, and control. In Zweben, M. and Fox, M. S., editors, *Intelligent Scheduling*. Morgan Kaufmann, San Francisco.

Warren, D. H. (1976). Generating conditional plans and programs. In *Proceedings of AISB76*, pages 344–354.