An intelligent decision support system for decision making under uncertainty in distributed reasoning frameworks

Tina Comes
Institute for Industrial Production
Karlsruhe Institute of Technology (KIT)
comes@kit.edu

Claudine Conrado
D-CIS Lab / Thales Research & Technology
Netherlands
claudine.conrado@d-cis.nl

Michael Hiete
Institute for Industrial Production
Karlsruhe Institute of Technology (KIT)
michael.hiete@kit.edu

Michiel Kamermans
D-CIS Lab / Thales Research & Technology
Netherlands
michiel.kamermans@d-cis.nl

Gregor Pavlin
D-CIS Lab / Thales Research & Technology
Netherlands
gregor.pavlin@d-cis.nl

Niek Wijngaards
D-CIS Lab / Thales Research & Technology
Netherlands
niek.wijngaards@d-cis.nl

ABSTRACT

This paper presents an intelligent system facilitating better-informed decision making under severe uncertainty as found in emergency management. The construction of decision-relevant scenarios, being coherent and plausible descriptions of a situation and its future development, is used as a rationale for collecting, organizing, filtering and processing information for decision making. The development of scenarios is geared to assessing decision alternatives, thus avoiding time-consuming analysis and processing of irrelevant information.

The scenarios are constructed in a distributed setting allowing for a flexible adaptation of reasoning (principles and processes) to the problem at hand and the information available. This approach ensures that each decision can be founded on a coherent set of scenarios, which was constructed using the best expertise available within a limited timeframe. Our theoretical framework is demonstrated in a distributed decision support system by orchestrating both automated systems and human experts into workflows tailored to each specific problem.

Keywords

Multi-criteria decision analysis (MCDA), scenario-based reasoning (SBR), distributed reasoning, multi-agent systems, emergency management.

INTRODUCTION

Decision making in emergency management presents all experts and decision makers involved with demanding challenges. Firstly, an emergency confronts society, economy and environment with substantial consequences. Secondly, the situation is very complex: the preferences of numerous actors regarding multiple objectives – and the tradeoffs involved – need to be respected. Thirdly, the information used as a basis for making a decision is prone to be uncertain: it is frequently not (yet) confirmed, noisy, lacking or even contradictory.

To deal with uncertainties, *scenarios* – being internally fully plausible, coherent and consistent descriptions of a situation and its future development – can be employed (Schnaars 1987). Scenarios facilitate reasoning under uncertainty: they support making robust decisions, i.e., the selection of an alternative, which performs sufficiently well under a variety of possible developments (Harries 2003). The use of scenarios challenges existing mental frames and avoids the cognitive biases in the estimation of probabilities (Wright and Goodwin 2009). To support decision making taking into account multiple goals and scenarios, techniques from Multi-Criteria Decision Analysis (MCDA) can be used (Comes, Hiete Wijngaards and Kempen, 2009). We now

Reviewing Statement: This paper has been fully double blind peer reviewed.

present a new method exploiting the integration of scenarios and MCDA in a distributed decision support system.

The first step in the scenario construction is the diagnosis, i.e., the collection of information regarding the emergency. Simply gathering information about the incident, the (ad-hoc) emergency management (including mitigation measures implemented so far, staff and equipment available), as well as the objectives and plans, is not sufficient in emergency management. Collecting and sharing information in an unstructured way has two major risks: first, there is the risk of information overload (too much redundant or irrelevant information is passed on), and second, it is not certain whether all relevant information will find its recipients.

Although in medium and longer term emergency management, decisions do not need to be made immediately, time is still restricted. Time criticality reduces the possibility of bringing together all experts involved in the scenario construction in person as is usually done in scenario planning, a discursive technique (Schnaars 1987). Furthermore, some experts may contribute to several decision problems and/or fulfil other tasks arising during the emergency. Experts may not have time to participate in the entire scenario construction process. Classical expert systems attempt to deal with this problem solving the decision problem autonomously, without interference of the users, by using a (limited) model of the domain and a set of data or assumptions (Turban and Watkins 1986). However, by excluding human experts, these systems require a vast, continuously updated, knowledge base covering all aspects of the available decision-options and all possible future developments (Dugdale 1996). In emergency management, this is clearly infeasible as emergencies can be characterised as rare, often unexpected events. Therefore, they defy overly standardized descriptions.

To facilitate medium and long(er) term decision making in emergency management, this paper presents an intelligent system supporting collaborative processing of information taking into account the bounded availability of experts as well as the inadequacy of standardized descriptions and evaluation of decision alternatives. Our approach combines the (cognitive) capabilities of multiple human experts and automated reasoning processes, each contributing specific expertise and processing resources to construct scenarios. To organize and structure information processing and sharing, *causal maps* (CMs) are used (Montibeller and Belton 2006). Finally, we show how the scenarios are used to evaluate each decision alternative with respect to multiple objectives.

The paper is structured as follows: in the first section, we explain how CMs are developed and how they link locally available expertise. The second section describes the construction of scenarios, including workflows that allow for the construction of plausible, coherent, consistent and decision-relevant scenarios. Thereafter, we show how this new approach facilitates decision making under severe uncertainty (i.e., when the situation defies quantitative descriptions). An example developed together with experts and users from the Danish Emergency Management Agency (DEMA) highlights the main features of our approach. Finally, we discuss the main aspects and give an outline of open questions.

DEVELOPING THE DECISION MODEL

Multi-criteria decision analysis has been frequently proven useful in long(er) term emergency management as it facilitates decision making in complex situations with respect to a variety of objectives (Papamichail and French 2005). In emergency management, the decision-making task is usually modelled as a choice among a number of alternatives based on a number of goals, making multi-attribute decision making (MADM) our preferred technique (Belton and Stewart 2002). A hierarchically structured attribute tree allowing for the evaluation of decision alternatives is elicited from the decision makers (Keeney and Raiffa 1976). This structure shows how strategic overall-objectives *(criteria)* are broken down first into criteria and finally into measurable attributes, taking the problem’s framing from an initially vague and intuitive understanding to a more formal description that can be analyzed mathematically. Despite their advantages in reducing complexity and arriving at a common understanding of the problem, both multi-attribute value and utility theory have some drawbacks when applied under severe uncertainty. While the first is a deterministic technique, the latter relies on probability distributions that are hard to define in emergency situations, as these are rare events (Ben-Haim 2000).

Scenario analysis is a well-established method for reasoning under severe uncertainty, as scenarios offer the possibility to consider several situation developments – regardless of their likelihood (Bunn and Salo 1993). Hereby it is possible to overcome cognitive biases such as overconfidence and to integrate fundamental risks that may be of very little probability into the reasoning framework (Schoemaker 1993). In our work, a scenario describes the current state of the situation and its development in the future.

To ensure its acceptance, each scenario should fulfil three conditions (Heugens and van Oosterhout 2001): *plausibility* (not going beyond the realm of possibility), *coherence* (having explicit logical connections explaining *how* the system evolves) and *consistency* (having no contradiction between its parts). The task of constructing scenarios that respect these conditions can be difficult. In scenario planning, scenarios are developed in a discursive procedure (Schnaars 1987), while formative scenario analysis (FSA) starts with identifying impact factors that influence the development of the situation (Scholz and Tietje 2002). Both approaches do not explicitly include interdependencies *(how* variables are related to one another), thereby requiring substantial effort to ensure coherence and consistency. In short- and even medium- to longer-term emergency management, the applicability of these techniques is limited, as time is usually restricted.

In this section, we present a novel approach to construct scenarios based on CMs that ensures plausibility, coherence and consistency as far as reasonable given bounded availability of experts and time. Furthermore, we show that our approach facilitates distributed reasoning by implementing the underlying workflows within a distributed approach called Dynamic Process Integration Framework (DPIF).

Decision model structuring using Causal Maps

The structuring of the decision problem is often characterised as one of the hardest, yet most crucial parts in developing decision support systems (Belton and Stewart 2002). Originally, CMs were developed as a problem structuring technique representing interlinked variables in a network (Montibeller and Belton 2006). Variables, which symbolize a feature characterising the situation, are depicted as nodes. Directed arcs describing cause-effect links connect nodes. The causality ensures that any CM is chronologically ordered: if a node i precedes a node j, the state of i influences the state of j causally. This implies that the state of i at time t influences the state of j at $t+\Delta$, $\Delta>0$. The temporal structure of the CM allows for the elimination of loops. Choosing the time steps appropriately, the CM can always be represented as a directed, acyclic graph.

A CM brings together several levels of a problem description. Considering only arcs and nodes reveals the structure of the problem. This *structural level* shows the variables that have an impact on the decision and their interrelations. At the *functional level*, the way in which a variable influences its successors is analysed in more detail. In emergencies, these relations are often uncertain and captured in a framework allowing for reasoning under uncertainty, e.g. in Bayesian Networks using probability distributions (Russell and Norvig 2003) or in Fuzzy Cognitive maps relying on Fuzzy Logic (Peña, Sossa and Gutiérrez 2008). The map's *informational level* captures the *results* associated to one instantiation of the system (Diehl and Haimes 2004), which can in our system be numerical, but also maps, recorded speech or text – according to what best fits the users' needs.

Distributed problem structuring using the Dynamic Process Integration Framework (DPIF)

In structuring the decision model, the first step in this framework is the elicitation of the MAVT attribute tree, where the decision makers define their objectives in terms of criteria and attributes (Keeney and Raiffa 1976), see Figure 1. The attributes – on the lowest hierarchical level of the attribute tree allow for *measuring* (or quantitatively estimating) the impact of the implementation of decision alternatives with respect to several objectives (Belton and Stewart 2002). Attributes correspond to states in the physical world as do the variables in a CM. On this basis we use the attributes as the intersection of the description of the situation and its development (represented in the CM) and the evaluation (represented in the attribute tree). The use of attributes facilitates filtering the variables in the CM: *relevant* variables (that need to be integrated) have an impact on at least one attribute.

To discover the relevant variables and their dependencies (i.e. nodes and arcs in the CM), a distributed approach based on the resolution of task dependencies is used. The DPIF lets experts (humans or automated systems) define their (reasoning) capabilities in terms of a task they can perform *(service)* and in terms of information this task requires. This reflects that an expert's output may rely on input he cannot determine autonomously. Our system connects experts via *software agents* that are an interface between the expert (human or automated) and the service-based discovery architecture the DPIF provides (Pavlin, Kamermans and Scafes 2009). Although it is necessary that the experts specify the *type* of information they need and can provide, the system allows for a flexible adaption of the formats used (e.g., images, spoken text or numbers).

For our system, we make the following assumption: Experts are considered as referring implicitly (in the case of humans) or explicitly (in the case of automated expert systems) to a local causal model that allows them to provide their service. Therefore, their reasoning processes are represented as local CMs (cf. Figure 2). These allow for mapping input from other experts to their own output

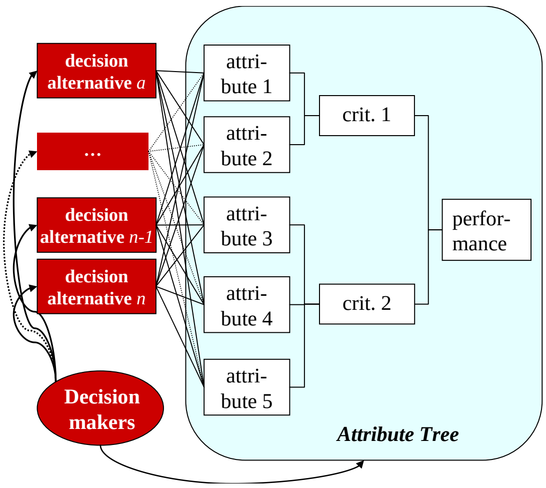

Figure 1: Attribute Tree

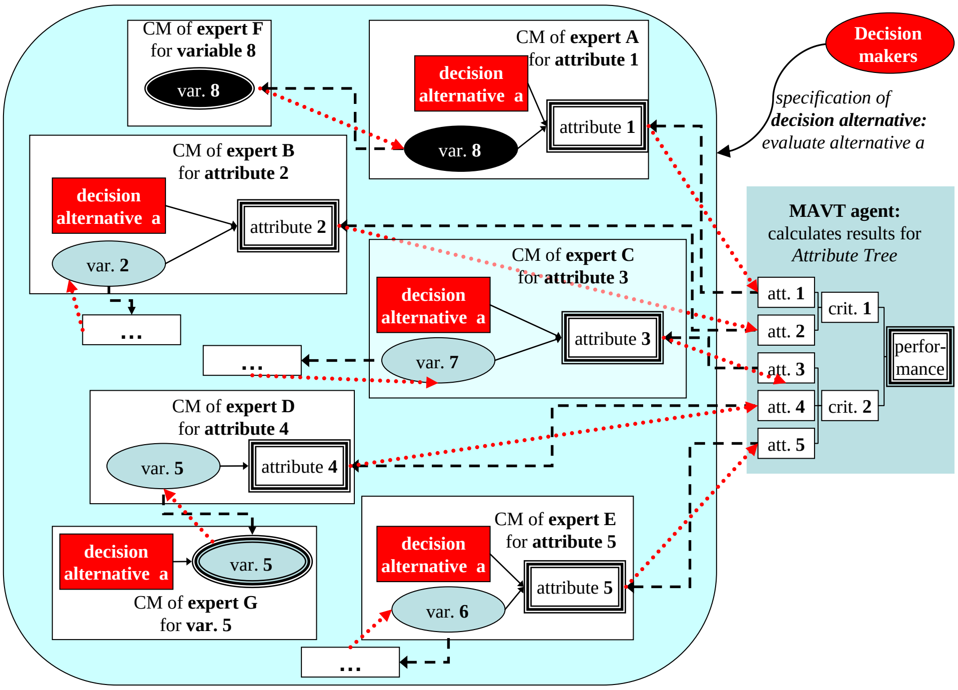

Figure 2: Constructing the global CM from local CMs

(captured in *local sink nodes*[1], see nodes with double borders in Figure 2). By iteratively combining the relevant local CMs, a global CM is developed in a top-down way (i.e., in the direction opposite to causation). In Figure 2, the black dashed arcs show how the local CMs are combined at configuration time, whereas the red dotted arcs correspond to the flow of information during computation. Strictly speaking, a system of collaborating experts arises combining their local knowledge such that the resulting system corresponds to a global CM.

When a decision must be made, the configuration process starts with a request from the decision maker to perform an evaluation for a set of decision alternatives (in Figure 2 only shown for one alternative *a).* The request for attributes' scores *given each decision alternative* initiates the construction of the CM. In the first step, DPIF agents look for experts that can provide information about the attributes. When such experts are found, they refer to their local CMs and indicate that in order to supply information about the attribute they depend on further information (see Figure 2).

Resolving the first information requirement described above does not finish the map elicitation process. Figure 2 shows that all experts *A* to *E* need additional information to provide their services (namely, information about the variables *2* and *5–8).* Again, the DPIF manages the requests for information. The system iteratively configures and expands a distributed reasoning model, which corresponds to the global CM This process finishes when the map covers all expertise needed to perform the scenario construction. This means that

- all dependent nodes are *sufficiently connected*, i.e. for all local CMs where the sink node has predecessor nodes, experts capable of providing the required information are identified;
- nodes without any predecessors are *independent* (depicted in black in Figure 2 e.g., var. *8*), i.e. the respective experts can determine the state of the corresponding variables autonomously (e.g., via measurements).

All local CMs have a single local sink node, which captures the information the respective experts are going to provide. The complete set of interconnected local CMs represents the *functional level* of the global CM, as it shows the expertise used to determine all relevant variables' states. To develop the corresponding CM on the *structural level*, the local CMs are merged. Individual experts are offered information which they defined relevant for fulfilling their task. Any expert is not confronted with other information that is processed in the global CM. This mechanism ensures reduction of the problem of information overload.

This system allows for flexible reaction to the problem at hand, as the processing of information (i.e. the *reasoning)* is *not* standardized (Pavlin, Kamermans and Scafes 2009) While the DPIF connects experts based on service descriptions, each expert is free to choose the manner of reasoning (e.g., algorithms, heuristics, best practices). This property is of great importance in highly varying, dynamic and unpredictable situations such as

[1] A sink node of a network graph is a node without successors. While locally, the output of each expert is captured in the sink node of its local CM, the sink node of the complete global CM merged with the attribute tree is the performance node.

emergencies. Particularly, it is possible that the expert adapts to the type and quality of information available (e.g. at first, information is uncertain, later it is confirmed). Yet, the structure of each expert's local CM (from the perspective of our model) usually remains unchanged in the course of the decision making process.

The global CM shows how to go *structurally* (in the direction of causation) from all independent nodes to the attributes required for finally evaluating the alternatives. Contrarily to expert systems, the CM does not encode the knowledge how this is achieved *functionally*. Rather, it specifies the *expert* responsible for determining the state of a variable and prescribes which input information is required. Several types of reasoning principles are hereby integrated into one map which can adapt their use flexibly to the problem at hand.

CONSTRUCTING SCENARIOS USING CAUSAL MAPS

This section shows how we use CMs to construct decision-relevant scenarios. Using the terms developed in the previous section, it is possible to clarify the term scenario further: a scenario is one complete instantiation of the CM (on the numerical level) together with its evaluation.

Scenario construction

The scenario construction in this framework follows an iterative procedure: starting with decision makers' request for attribute scores given an alternative, all experts linked in the CM are activated and asked to provide their specific service (configuration in a top-down way, cf. Figure 2). Then, the information is processed in a bottom-up manner following the links in the CM. Each expert uses his local knowledge and procedures (represented as local CMs) to determine the possible state(s) (multiple in case of uncertainty, see below) of the variable for which he agreed to provide results given the information he receives.

The processing of information starts with the analysis of *independent nodes* in the map (depicted in black, e.g. var. *8* in Figure 3). The experts providing information about the independent nodes assess the states of the corresponding variables (based on their local knowledge).

If there is uncertainty about the state of a variable, an expert can pass on several possible estimates for one variable. This information can be encoded as a set of numerical values, a number of maps, some text files etc. All output variants must fit to the input that was used to determine these estimates. Figure 3 shows a part of the workflow where a set of scenarios arises. Assume that variable 8 is prone to uncertainty. The responsible expert F decides to transfer three possible states *(I, II* and *III).* Then, expert A assessing attribute *1* must determine this attribute *under all states of* variable *8*. Expert *A* himself can determine a set of possible and relevant states for each of the states of variable *8*. In the example shown in Figure 3, *A* passes on multiple possible states for each of the three possible states of variable *8*, namely, two states for *8=I* and *8=III*, and three states for *8=II*. The result is a total number of seven partial scenarios[2] for attribute *1* under decision alternative *a*.

In brief, whenever there is uncertainty on how the 'story' develops, the partial scenarios bifurcate in a number of possible 'storylines'. The arising scenarios can therefore be understood as a way of expressing uncertainty reflected in a range of possible and relevant states for each variable. In this manner, scenarios are completed iteratively – starting at the independent nodes and following the CM until the attributes are reached – by the development of consistent conditional partial scenarios. The number of variable states passed on must be chosen carefully to avoid the problem of combinatorial explosion (see discussion section).

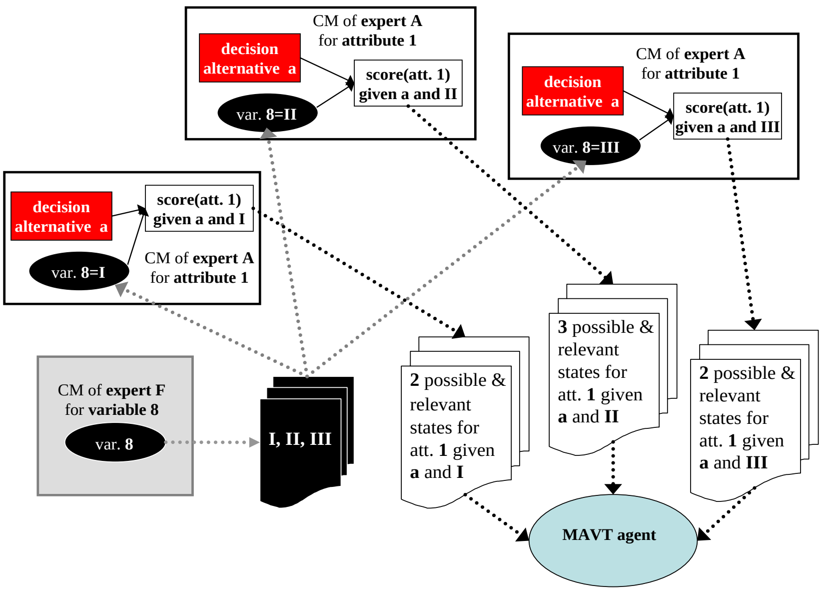

Figure 3: Scenario Construction

This approach ensures the

[2] The scenarios are called *partial* here as a complete scenario encompasses a description of *all* the variables in the CM, not only those relevant for the determination of just one attribute.

consistency of each scenario from the beginning onwards, under the assumption that the local assessments of the experts reflect the "state of the art" as good as possible, given their current knowledge. Contrarily to Formative Scenario Analysis (FSA) – where first a vast set of scenarios is created by combining all possible states of all variables arbitrarily, which is afterwards reduced by applying a pair-wise consistency assessment carried out by experts (Scholz and Tietje 2002) – in our approach consistency is ensured in a distributed manner by the experts providing each piece of information. In this way, the number of scenarios is kept smaller than in FSA. Furthermore, the scenarios are *coherent*, as not only the states of all variables but also a description of the way in which they are interlinked is part of each scenario. Finally, the *plausibility* of each scenario is enhanced as our approach allows for understanding the underlying interdependencies between the variables, integrating the best available expertise to assess their state(s). The fulfillment of the three conditions above should therefore be understood under the underlying assumption that *the experts' local assessments reflect the best expertise that is currently exploitable for the decision problem at hand*.

Evaluation

Our approach is targeted at supporting decision makers in the evaluation of decision alternatives under several scenarios. To accomplish this aim, the CM is connected to the MAVT attribute tree (cf. Figure 4). For each decision alternative a_i $(i=1,\dots,n)$ an assessment of the attributes' scores is requested. Following the approach described previously, these scores are determined.

In case the states of all variables are deterministic, one set of attribute scores per decision alternative is assessed. In this case, the usual MAVT approach is followed (cf. Belton and Stewart 2002). If there is uncertainty, however, a set of scenarios $S(a_i)$ is created for each alternative a_i. The complete set of constructed scenarios is denoted Ω. When all attribute scores for all $s \in \Omega$ are determined, each scenario s is evaluated using the attribute tree and the decision makers' intra- and inter-criteria preferences, resulting in a performance $p(s)$ (Belton and Stewart 2002). This technique allows for *comparing* the performances of each decision alternative under a variety of scenarios.

To facilitate decision-making encompassing all scenarios, a further MAVT- step is performed (cf. Figure 4). To this purpose, weights $w_j[s_j(a_i)]$ reflecting the *relative importance* of each scenario $s_j(a_i)$ are elicited from the decision makers analogue to the usual preference elicitation, where $j=1,\dots,m$, $|S(a_i)|=m$ and $\sum_{j=1}^{m} w_j = 1$.

These weights to not reflect the probability or likelihood of a scenario, rather they are used to take into account the deviation of the evaluation (overall or with respect to several criteria) from a desired value. For each decision alternative a_i the performances of all scenarios $s_j(a_i)$ are aggregated as $p(a_i)=\sum_{j=1}^{m} w_j \cdot p(s_j(a_i))$ (cf. Figure 4). In this manner, *robust* decision making, i.e. the choice of an alternative that performs sufficiently well for a set of scenarios (Ben-Haim 2000) is supported. This approach can be particularly useful when Ω is large, or when there is overconfidence in a narrow range of scenarios, neglecting the significance of all other possible developments; a situation frequently encountered in emergency management (Wright and Goodwin 2009).

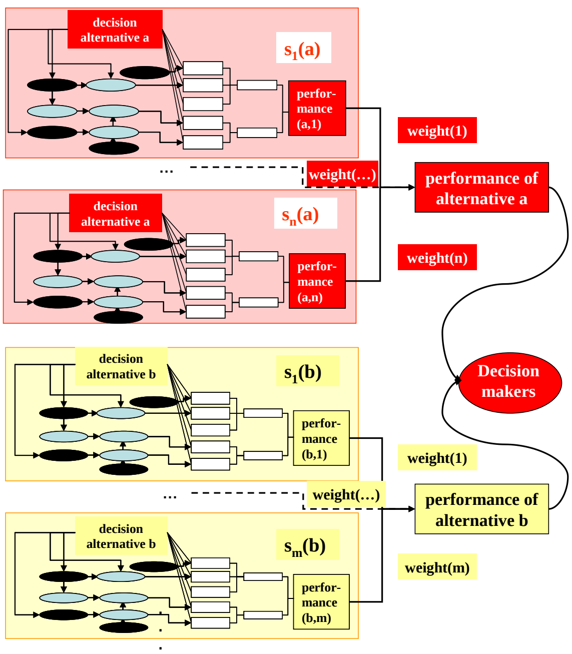

Figure 4: Robust decision making

To enhance acceptance and to facilitate consensus building on the preferences (e.g., as captured in scenario weights), it is not sufficient to provide the decision makers only with the total performance for each alternative. Rather, the result for the scenario should be presented in more detail. A selection criterion is needed as, unfortunately, the cognitive capacity of decision makers is limited (ca. five scenarios at a time, Godet 1990). The integration of CM and attribute tree facilitates the use of the evaluation of scenarios to select for each decision alternative the scenarios for detailed analysis. This approach ensures that the chosen scenarios are the most distinct with respect to the decision makers' actual assessment. For each set $S(a_i)$, the scenarios with the worst, the best and the performance closest to the median are selected and presented to the decision makers in detail, adapting the usual choice of a pessimistic, an optimistic and a baseline scenario (Schnaars 1987) to the MAVT framework. This approach allows making the spread of possible evaluations for each alternative visible and facilitates robust decision-making, when e.g., a minimum worst-case performance must be reached.

EMERGENCY MANAGEMENT EXAMPLE

The approach presented in the previous sections is illustrated by means of an emergency management example that has been developed together with the Danish Emergency Management Agency (DEMA).

Situation Description

Two freight trains crash at the central train station in Odense, Denmark, causing the leakage of chlorine from a ruptured tank wagon. The Hazmat unit – a specialised unit from the nearest DEMA rescue centre – covers the leak and stabilizes the situation temporarily. A permanent solution requires the chlorine to be transferred to another tank. This transfer is fraught with the risk of a further chlorine leakage creating a (lethal) plume over the downwind area. In order to protect the population, a decision must be made regarding which preventive measure should be applied: *evacuation* of downwind areas or *sheltering* in house.

Problem Structuring

To develop the decision model for the example, an attribute tree is elicited which includes the criteria *Health*, *Effort*, *Economic losses* and *Impact on society*. To illustrate our approach, we focus on *Health* topics and first show how the attributes *Number of ill in hospital to be sheltered* and *…evacuated* are determined.

The DPIF looks for experts whose service is providing scores for these attributes (e.g., a health expert). To determine the scores, the identified expert's indicate the need for further information on the number of ill in the hospitals that are (potentially) exposed and on the alternative implemented. This is depicted in Figure 5, where both nodes are connected to the attributes. Again, grey dashed arcs represent the way the nodes are connected during the configuration phase, whereas red dotted arcs indicate the flow of information.

The CM configuration process continues by determining the information required by the experts providing information about the *Decision alternative* and the *Number of ill in hospitals exposed*. For the former the information is provided by the decision makers, whereas for the latter an expert is consulted who requires knowledge on the hospitals exposed. This is represented in Figure 5 as a further link connecting *Number of ill in hospital exposed* to *Hospitals exposed*. Continuing this process iteratively, the global CM expands until all independent nodes (represented in black) are reached. Figure 5 represents the partial CM that is developed by merging the involved experts' local CMs.

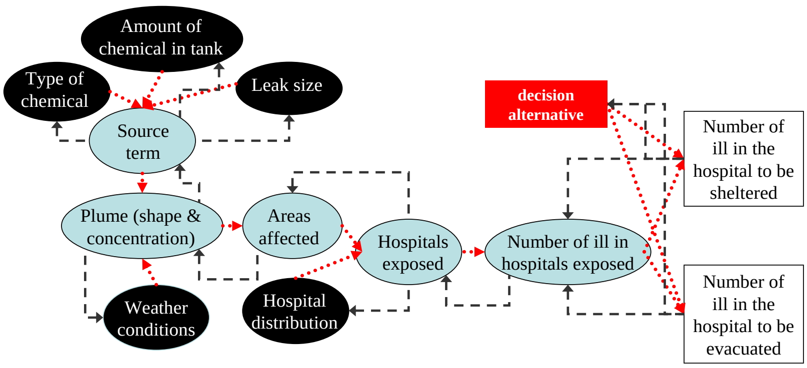

Figure 5: Partial CM to assess "Number of ill in hospital"

Scenario construction

We now illustrate how scenarios are constructed for each decision alternative by processing information bottom-up following the causal links in the CM to determine the score of each variable in the path. We refer once more to the partial causal map in Figure 5. First, the scores for the independent nodes are determined. Assume that the relevant expert (the Local Incident Commander, LIC) is uncertain about the amount of chemical left in the tank. Consequently, he decides to pass on not a single but three estimates, representing the cases that a *small* (10 t), a *medium* (30 t) or a *large* amount of chlorine (60 t) is left in the tank. A further expert (the Duty Hazmant Officer, DHO), who determines the source term, relies on this information. For each of the estimates the LIC passed on to him, the DHO develops the respective source term (see Figure 6, showing the LIC's local CM in the light blue box and the DHO's local CM in dark blue boxes). This three-fold uncertainty regarding the amount of chlorine is passed on further (via the CM) to all nodes depending on the amount of chlorine in the tank. When further uncertainty arises, this is accommodated by passing on the respective number of possible states for each of the partial scenarios described by the predecessor nodes.

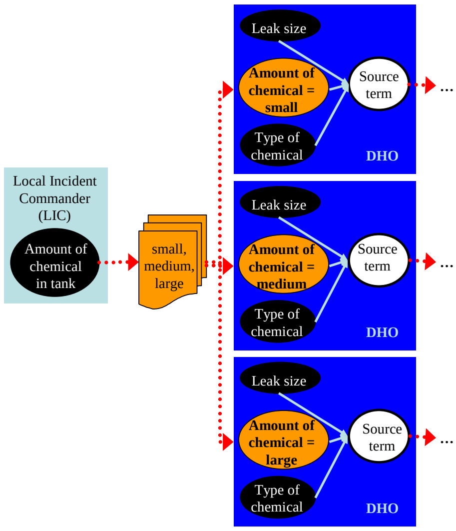

Figure 6: Start of scenario construction

When distinct branches of the CM contribute to the determination of one variable and each of these includes one uncertain variable, it is important that they can be combined in a manner that ensures consistency. If n uncertain branches coming together in one node are independent (i.e. all paths do not have any nodes in common), combinations of all possible states for all n variables must be considered. In Figure 7, e.g., the source term and the weather conditions are independent, so $2 \cdot 3 = 6$ partial scenarios are considered in the assessment of "Areas affected".

It is important to notice that the number of decision-relevant scenarios does not necessarily increase following the CM. Figure 7 shows, e.g. that there are six distinct affected areas to be considered. The number of scenarios passed on for *Hospitals exposed* is only two. This can happen if there is, e.g., only *one* hospital next to the incident location. In some of the scenarios, this hospital may be affected, while remaining unaffected in other scenarios. By filtering the scenarios according to their relevance for the decision at hand (measured by the attributes), the number of scenarios completed is kept smaller than in FSA.

Evaluation of decision alternatives

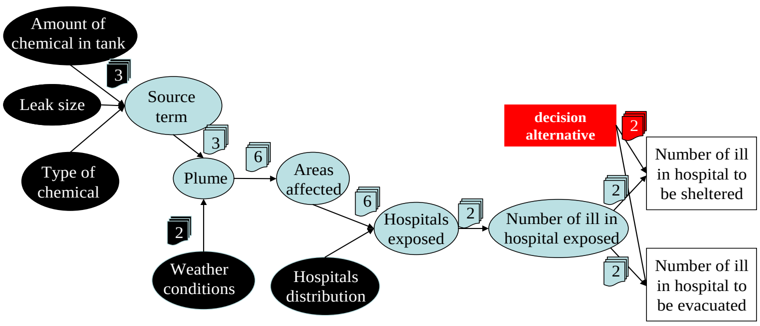

Figure 7: Multiplicity of partial scenarios in the example

When the CM is fully assessed for all decision alternatives with respect to *all* attributes, the attribute scores for each scenario are evaluated using MAVT techniques. To overcome cognitive biases such as overconfidence and to avoid that the decision makers focus on a limited set of particularly impressing scenarios, an additional aggregation step is performed. To this end, scenario weights reflecting the relative importance of each scenario are elicited from the decision makers. In this example, equal weights for the scenarios of each decision alternative were used. This results in a higher total performance for the sheltering alternative (0.66 vs. 0.59 for evacuation).

In addition to these results, the decision makers are provided with stacked bar charts showing the performance of the worst, medium and best scenarios for both decision alternatives (see Figure 8, reading from left to right). This chart facilitates the assessment of robustness. Figure 8 shows, e.g., that the worst scenario for alternative evacuation results in a much better performance than the worst scenario for sheltering. A sensitivity analysis, varying the scenario weights, provides further support (cf. Comes et. al., 2009).

CONCLUSION

We propose a framework facilitating medium and longer term decision-making under uncertainty by considering scenarios: different possible future developments of an emergency situation. Scenarios, being plausible, coherent and consistent situation descriptions are easily understandable and help overcome cognitive biases (Wright and Goodwin, 2009). This paper introduces a novel approach to construct scenarios in a distributed manner. Causal maps (CMs)

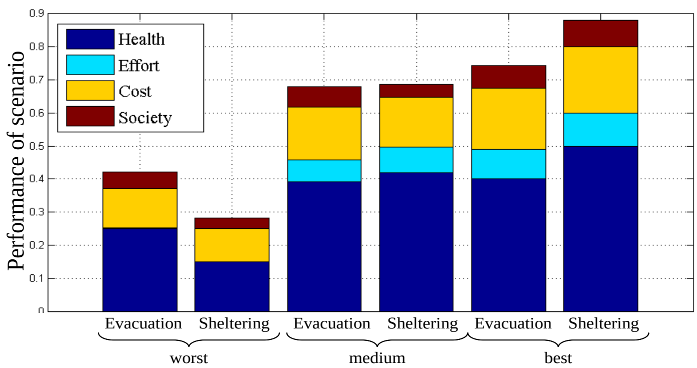

Figure 8: Evaluation of evacuation and sheltering

make it possible to take into account the interdependencies between all relevant variables. In this manner, the consistency and coherence of each scenario can be ensured to a reasonable level given that availability of experts is bounded and time for the decision at hand is not unlimited.

Using the MAVT attribute tree as starting point of the scenario construction ensures that only information relevant for the decision at hand needs to be processed. Furthermore, this reduces the number of relevant scenarios, as only scenarios that are distinguishable with respect to the attributes need to be considered.

Our approach facilitates distributed reasoning and involves human experts as well as automated systems. Each expert has the opportunity to choose the methods to determine the state of each variable freely and to adapt these as the situation evolves (e.g., when previously lacking or uncertain information becomes known). Altogether, our approach results in overall distributed heterogeneous problem solving behaviour, enabling the system to adapt to a dynamic, highly varying environment such as encountered in emergency management.

The sharing and processing of information is coordinated by having each expert specify which service can be provided and which information is relevant. Each expert's expertise is represented as a local CM with the additional benefit that the expert is provided with all information he judged necessary for his task. Simultaneously, an expert is *not* confronted with information irrelevant for performing a particular task, thus reducing information overload. As the scenarios developed in our framework are established by (human or artificial) experts, a major drawback in the use of computer-based systems, namely the problem of missing acceptance and trust in anonymous systems (Engelmann and Fiedrich 2009), can be circumvented.

To come to a robust recommendation, a twofold approach is proposed. First, a detailed analysis of the scenarios with the worst, the best and a medium performance for each decision alternative is provided. Second, an aggregated overall performance, which encompasses the evaluations of all scenarios for each decision alternative, allows for overcoming overconfidence in a small range of scenarios. Thus, decision makers can gain deeper insights into the decision situation than those provided by standard methods from both SBR and MCDA.

Our approach and its application during environmental emergency management is further researched in close collaboration with emergency management authorities, e.g. with the Danish Emergency Management Agency (DEMA). Our research aims at tailoring this approach so that it fits best the decision makers' and experts' needs. To this end, a validation experiment is planned early 2010 involving a multi-user demonstrator in which both human expertise and automated systems are involved. The experiment's objectives include an analysis of how much individual (human) experts need to know about the global CM and the extent to which they are able to work on multiple input variants, given a limited timeframe.

A number of open questions arise from our approach. To accelerate the processing of information in case of an emergency, the local CMs from all (potentially) involved experts and the conditions under which they are willing and able to provide their service need to be elicited. To this end, we collaborate closely with Dutch and Danish emergency management agencies. The DPIF negotiation algorithms, ensuring that the best experts capable of providing information within a certain timeframe are found, must be further developed to include methods to identify and manage missing expertise (i.e., required input information that can *not* be delivered by any expert). So far, in each example that we analysed, the number of scenarios constructed has been manageable. Yet, in highly uncertain contexts the number of scenarios may become too large (due to combinatorial effects) to allow for timely processing. Our current research includes investigating a scenario pruning or prioritisation mechanism. Another issue is the question of how multiple decision making processes (performed in parallel on different hierarchical levels or sequentially) can be coordinated such that an optimal set of decisions can be identified between these decision making processes. Finally, the results must be presented in a transparent and easily understandable way. Visualisations of (partial) scenarios for decision making remains of importance: a number of presentation alternatives are investigated. Additionally, a documentation of uncertainties related to each scenario must be developed and adequately presented.

ACKNOWLEDGMENTS

This research has been conducted in the Diadem project, funded by the European Union under the Information and Communication Technologies (ICT) theme of the 7th Framework Programme for R&D, ref. no: 224318, www.ist-diadem.eu. The authors wish to thank their partners and the ISCRAM 2010 reviewers for the valuable comments.

REFERENCES

1. Belton, V. and Stewart, T. (2002) Multiple Criteria Decision Analysis. An Integrated Approach, Kluwer Academic Publishers, Boston
2. Ben-Haim, Y. (2000) Robust rationality and decisions under severe uncertainty, *Journal of the Franklin*

Institute 337, 2-3, 171-199.

3. Bunn, D. and Salo, A. (1993) Forecasting with scenarios, *European Journal of Operational Research* 68, 3, 291-303.
4. Comes, T., Hiete, M., Wijngaards, N. and Kempen, M. (2009) Integrating Scenario-Based Reasoning into Multi-Criteria Decision Analysis, in: J. Landgren und S. Jul. (eds.) *Proceedings of the 6th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2009)*, Gothenburg.
5. Diehl, M. and. Haimes, Y. (2004) Influence diagrams with multiple objectives and tradeoff analysis, *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, 34, 3, 293-304.
6. Dugdale, J. (1996) A cooperative problem-solver for investment management, *International Journal of Information Management* 16, 2, 133-147.
7. Engelmann, H. and Fiedrich, F. (2009) DMT-EOC – A combined system for the Decision Support and Training of EOC Members, in: J. Landgren und S. Jul. (eds.) *Proceedings of the 6th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2009)*, Gothenburg.
8. Godet, M. (1990) Integration of scenarios and strategic management: Using relevant, consistent and likely scenarios, *Futures*, 22, 7, 730-739.
9. Harries, C. (2003) Correspondence to what? Coherence to what? What is good scenario-based decision making?, *Technological Forecasting and Social Change*, 70, 8, 797-817.
10. Heugens, P. and van Oosterhout, J. (2001) To boldly go where no man has gone before: integrating cognitive and physical features in scenario studies, *Futures*, 33, 10, 861-872.
11. Keeney, R. and Raiffa, H. (1976) *Decisions with Multiple Objectives*, John Wiley & Sons, New York.
12. Montibeller, G. and Belton, V. (2006) Causal maps and the evaluation of decision options - a review, *Journal of the Operational Research Society*, 57, 779-791.
13. Papamichail, K.N. and French, S. (2005) Design and evaluation of an intelligent decision support system for nuclear emergencies, *Decision Support Systems*, 41, 1, 84-111.
14. Pavlin, G., Kamermans, M. and Scafes, M. (2009) Dynamic Process Integration Framework: Toward Efficient Information Processing in Complex Distributed Systems, In *Proceedings of IDC 2009 - 3rd International Symposium on Intelligent Distributed Computing, October 13-14, 2009, Ayia Napa, Cyprus.*
15. Peña, A., Sossa, H. and Gutiérrez, A. (2008) Causal knowledge and reasoning by cognitive maps: Pursuing a holistic approach, *Expert Systems with Applications*, 35, 1-2, 2-18.
16. Russell, S. and Norvig, P. (2003) *Artificial intelligence*, Prentice Hall, Upper Saddle River, NJ.
17. Schnaars, S. (1987) How to develop and use scenarios, *Long Range Planning* 20, 1, 105-114.
18. Schoemaker, P. (1993) Multiple scenario development: Its conceptual and behavioral foundation, *Strategic Management Journal*, 14, 3, 193-213.
19. Scholz, R. and Tietje, O. (eds.) (2002) *Embedded Case Study Methods. Integrating Quantitative and Qualitative Knowledge* Sage Publications, Thousand Oaks.
20. Turban, E. and Watkins, P. (1986) Integrating Expert Systems and Decision Support Systems, *MIS Quarterly*, 10, 2, 121-136.
21. Wright, G. and Goodwin, P. (2009) Decision making and planning under low levels of predictability: Enhancing the scenario method, *International Journal of Forecasting*, 25, 4, 813-825.