On the use of automated planning for crisis management

Loïc Bidoux
Mines Albi - Université de Toulouse
Thales Communications & Security
loic.bidoux@mines-albi.fr

Jean-Paul Pignon
Thales Communications & Security
jean-paul.pignon@thalesgroup.com

Frédérick Bénaben
Mines Albi - Université de Toulouse
frederick.benaben@mines-albi.fr

ABSTRACT

Automated planning is a domain of Artificial Intelligence which aims to study the deliberation process used to choose and organize actions by anticipating their expected outcomes. In this paper, we discuss the use of automated planning techniques in crisis management contexts. To begin with, the crisis management planning problem is formalized in light of the conceptual model for automated planning. In addition, we describe the conceptual scheme of an information system generating action plans in order to support decision-makers in crisis management. Finally, a proof of concept implementation of the aforementioned system is presented.

Keywords

Crisis management, automated planning, collaborative plan, multicriteria decision analysis, ChoPlan

INTRODUCTION

A *crisis* (or critical situation) can be defined as “a situation with long-term consequences due to an event that has caused extensive damages and losses resulting in an interruption of one or more critical activities within some part of the world” (CCA 2014). Such situations may, for instance, result from natural disasters (tsunamis, earthquakes, floods. . .), industrial accidents (e.g. an explosion of a production facility) or from malevolent acts (terrorist attacks for instance). Managing critical situations still represents a major challenge for our society. Indeed, between 2006 and 2015, crisis were responsible for approximately 70,000 deaths per year according to the last statistical review of the international Emergency Event Database (Glasser and Guha-Sapir 2016). Most of the practitioners (see for example (Altay and Green 2006)) describe the life cycle of crisis management using four phases: *prevention*, *preparation*, *response* and *recovery*. This study mainly focuses on the response phase.

When a crisis occurs, many organizations are mobilized to respond. The term *organization* is used here in its broadest meaning ; it may refer to state emergency services such as fire department or police as well as companies or associations. The management of the crisis response is generally charged to a so called *crisis cell* composed of various stakeholders: one or more decision-makers assisted by relevant actors and experts in light of the situation. The mission of the crisis managers is delicate since they have to coordinate the collective work of many organizations in a stressful and time-constrained environment. In most cases, involved organizations are both heterogeneous (at cultural, functional and technological levels) and poorly, if at all, trained to work together in the emergency context of the crisis. This inevitably leads to collaboration issues (imperfect definition of the goals to achieve, incomplete information sharing, poor coordination of actors. . .) that limit the efficiency of the response. These problems have been highlighted by many feedbacks from past crises (Van De Walle and Turoff 2007; Treurniet et al. 2012) showing that maturity of the coordination between stakeholders is a limiting factor in crisis management.

To address these collaboration issues, preventive contingency plans are generally prepared before critical situations arise. Unfortunately, as they are generated when the operational context of the collaboration is not entirely known and because real crisis situations may diverge rapidly from planned ones, these plans are often imperfect. Dwight D. Eisenhower has said in this regard that "Plans are worthless, but planning is everything" (Eisenhower 1957). Through this statement, the military strategist emphasize the necessity for a plan to fully fit into the operational context justifying its creation. In order to avoid the aforementioned limitations, this study only focuses on action plans that are produced once the crisis event has occurred.

One should note that the crisis cell operate at a strategic level. As a consequence, *collaborative plans* produced by its members are singular in some regards. They involve many organizations and focus less on the precise description of the actions to execute (which belong in the expertise area of the responders) rather than on the orchestration of these actions. Indeed, according to the principle of subsidiarity, decision-makers can order to involved partners what to do and when to do it but they never advise them on the way these tasks should be performed. Besides, planning activities performed by the members of the crisis cell are the only ones to be considered in this study. Other planning approaches for crisis management (for instance, planning to optimize victim evacuation by taking account of the road traffic (Aligne and Savéant 2011)), although independently interesting, are considered as complementary to our approach but are not discussed further.

Automated planning is a domain of Artificial Intelligence which studies the deliberation process that chooses and organizes a set of actions according to their expected outcomes. In this paper, we discuss the use of automated planning techniques in order to support the planning activities of the crisis cell's members. First, a brief overview of automated planning and its conceptual model is given. The assumptions of this conceptual model are then analyzed in the context of critical situations as we present what we believe is the correct way to model the crisis management planning problem. This problem encompass two distinct problematics amongst which only the first one is studied. Indeed, the third section focuses on this problematic by describing the core features that an information system generating plans in order to support the crisis managers should possess. Next, a proof of concept implementation of such a system is presented. Finally, a concluding section sums up these results and suggests some perspectives.

AUTOMATED PLANNING OVERVIEW

Automated planning addresses the problem of finding a sequence of actions to execute on a system Σ in order to achieve a specified goal state from a given initial state. In its broadest meaning, the planning activity encompasses both the plan elaboration as well as its execution. A conceptual model for planning has been proposed in the reference textbook *Automated Planning*: *Theory and Practice* (Ghallab et al. 2004). It is based on the interactions of three components: a planner, a controller and a system Σ (cf. figure 1).

1. The *planner* generates an action plan based on the description of the system Σ, its initial state and the objectives to achieve ;
2. The *controller* executes the plan generated by the planner thereby producing actions on the system. To do so, it relies on information (possibly incomplete) from the current state of the system ;
3. The *system* Σ evolves in response to actions performed by the controller or external events that may occur.

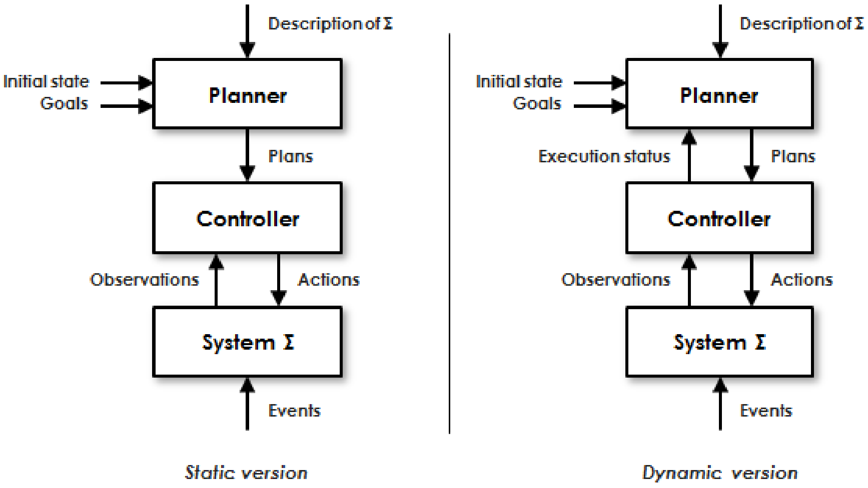

Figure 1. Conceptual model for planning (Ghallab et al. 2004)

The controller is usually assumed to be robust enough to handle the differences that may exist between the real world and its model. However, if this assumption is not acceptable, the controller may return an execution status to the planner allowing the latter to produce a new plan whenever the observed situation and the expected one diverge. In this case, the term *dynamic planning* is used to denote that the construction of the plan and its execution become closely linked (cf. feedback loop on the right part of figure 1).

Since planning is concerned with choosing and organizing actions in order to modify the state of a system, a model to specify and describe the dynamic evolution of systems is required. The *state-transition system* model (Dean and Wellman 1991) is generally used to this end. Formally, a state-transition system is a 4-tuple $\Sigma = (S, A, E, \gamma)$, where :

- S is a finite or recursively enumerable set of states ;
- A is a finite or recursively enumerable set of actions ;
- E is a finite or recursively enumerable set of events ;
- $\gamma : S \times A \times E \rightarrow \mathfrak{P}(S)$ is a state-transition function (with $\mathfrak{P}(S)$ denoting the power set of S).

Modelling Σ as a state-transition system is equivalent to describing it by its state $s \in S$. As a consequence, an evolution of the system Σ corresponds to a state transition $s \rightarrow s'$. In the conceptual model for planning, the transition from a state s to a state s' is represented by the execution of an action/event pair (a, e) in s. The neutral action α and the neutral event ϵ are introduced in order to describe transitions caused only by actions or caused only by events. In such cases, notations $\gamma(s, a)$ and $\gamma(s, e)$ are used in place of $\gamma(s, a, \epsilon)$ and $\gamma(s, \alpha, e)$. One should note that although actions and events both contribute to the system evolution, their semantics differ. Actions are controlled by the people in charge of the plan execution whereas events are transitions that are beyond their control. The latter can be caused by the internal dynamic of the system or the evolution of its environment.

The function γ indicates the state s' towards which the system Σ may evolve from state s in response to an action a, an event e or an action/event pair (a, e). Thus, the size of the set returned by $\gamma(s, a)$, which is denoted as $|\gamma(s, a)|$, provides several information regarding the execution of a in s. If $|\gamma(s, a)| > 0$, the action a is said to be applicable in s since the system Σ can transition towards at least one state s' whenever a is executed in s. In addition, if $|\gamma(s, a)| \leq 1$, the execution of a in s is said to be deterministic. Indeed, if the action a is executed in s, the system Σ can only transition to a single state s'.

Furthermore, the conceptual model for planning (Ghallab et al. 2004) introduces eight assumptions that are used to characterize the different classes of planning problems.

Assumption A1 (Finite Σ). The system Σ has a finite set of states.

Assumption A2 (Fully Observable Σ). The state of the system Σ is fully observable. Therefore, the observations used by the controller are all perfect.

Assumption A3 (Deterministic Σ). The system Σ is deterministic if for every state s, every action a and every event e, $|\gamma(s, a, e)| \leq 1$. Thus, in response to an action/event pair (a, e) applicable in s, the system Σ can only transition towards a single state s'.

Assumption A4 (Static Σ). The system Σ has no internal dynamics therefore the set of events E is empty.

Assumption A5 (Restricted Goals). One can only specify goals to achieve as properties of the final state of the system. As a consequence, a solution to a planning problem is any sequence of state transitions which leads to a final state $s \in S_G$ with S_G the set of states in which the objectives are verified.

Assumption A6 (Sequential Plans). Solution plans are represented by a finite sequence of linearly ordered actions thus there is no parallelism involved in them.

Assumption A7 (Implicit Time). Actions and events have no intrinsic duration and are considered as instantaneous state transitions.

Assumption A8 (Offline Planning). Dynamic planning mechanisms are not required.

The term *classical planning* refers to the class of problems obtained when all the assumptions of the conceptual model are considered simultaneously. Studying the classical planning problem is fundamental because most of the planning problems are defined by extending the latter ; like the crisis management planning problem presented in the next section for instance.

Classical Planning Problem (Ghallab et al. 2004). Given $\Sigma = (S, A, \gamma)$ a state-transition system, $s_0 \in S$ an initial state and $S_G \subseteq S$ the set of the states verifying the restricted goals G, find a sequence of actions $\langle a_1, \cdots, a_n \rangle$ corresponding to a sequence of states $\langle s_0, \cdots, s_n \rangle$ such that $s_1 = \gamma(s_0, a_1), \cdots, s_n = \gamma(s_{n-1}, a_n)$ and $s_n \in S_G$.

As the assumptions of the conceptual model are rather restrictive, some real world problems can not be modelled accurately by the classical planning problem. In such cases, one may consider other classes of problems by relaxing some of the aforementioned assumptions. These new classes of problems are harder to solve and constitute various subdomains of automated planning. For instance, the assumption A7 is relaxed in temporal planning while assumptions A1 and A5 are not considered in preference-based planning.

In the following section, assumptions of the conceptual model are analyzed in the context of crisis management. Doing so, we present what we believe is the correct way to model the crisis management planning problem.

AUTOMATED PLANNING FOR CRISIS MANAGEMENT

As mentioned in the introduction, dynamics of critical situations may be relatively unpredictable. As a consequence, assumption A8 has to be relaxed in the context of crisis management. Using a dynamic planning approach is more complex as it implies to monitor the crisis response in order to detect potential divergences between the actual situation and the expected one. If such a case occurs, the current plan must either be adapted or completely redesigned to take into account the novelty of the situation. On the other hand, other assumptions may be considered thanks to dynamic planning thus offering an interesting trade-off between the complexity of the planning problem to solve and the need to adjust the produced plan during its execution. Indeed, assumptions A2, A3 and A4, even if not unconditionally true in a crisis management context, may be considered as valid hypotheses given the possibility to restart the planning procedure if needed. Doing so, the system may be considered fully observable, deterministic and static. Nevertheless, replanning mechanisms will be triggered if a (previously unknown) relevant information becomes available, a responder action does not produce its expected outcome or if an external event significantly impacts the system environment.

Crisis management planning problems are rather different from traditional automated planning problems. They tend to be easier to solve than the latter as they are usually less combinatorial while being much harder to model. Indeed, it is quite difficult to represent the goals to achieve as determining the best strategy to efficiently handle the situation is generally crisis-specific and might be subject to debate amongst decision-makers. It results from these considerations that assumption A5 is too restrictive in a crisis management context. In order to fully capture decision-makers needs, extended goals and preferences (Gerevini and Long 2006) should be preferred in place of restricted goals. Goals are strong constraints that must be verified by the solutions of the problem whereas preferences are soft constraints that may not be verified in a solution plan. Extended goals encompass both *final state goals* (which correspond exactly to the restricted goals of assumption A5) and *trajectory goals* that formalize requirements on the sequences of actions characterizing the solution plans. For instance, they can be used to forbid transitions to specific states or to force the system to be in a particular state before another one. Similarly to objectives, preferences can represent conditions on the final state (*final state preferences*) or intermediate states (*trajectory preferences*) of the solution plans. In addition, they can also refer to any numeric variable (*numerical preferences*) defined in the problem. This gives them a rich expressive power since they can describe costs and risks to be minimized (logistic cost, response time. . .) as well as quantities to be maximized (some quality levels for example). When numerical preferences are used to describe the state of the system, it becomes impossible to consider assumption A1. Relaxing this assumption is mainly a technicality ; although it may affect employed automated planning techniques, it has no operational impact. Furthermore, in a crisis management context, one should consider the approach described in (Bidoux 2016) which use the Multi-Attribute Utility Theory formalism (Dyer 2005; Grabisch and Labreuche 2010) along with a Choquet integral (Choquet 1953) to represent the aforementioned preferences. It enforces commensurability between the considered preferences (trajectory and numerical preferences are not by default and are relatively hard to compare) which greatly simplifies both the modelling of the planning problem as well as the comparison of candidate solution plans. Moreover, this approach offers additional expressive power advantages as it allows decision-makers to aggregate several preferences together and to take into account interactions that may exist between the preferences.

Time is often a critical resource in a crisis management setting, where parallelizing responders actions whenever possible in order to optimize response efficiency is generally very important or even crucial. This indicates that assumption A6 of the conceptual model for planning is not suitable in such contexts. Parallelization can be performed either while planning (increasing the problem difficulty) or as a post-treatment once a solution plan is found. The latter approach may produce potentially suboptimal plans from a scheduling point of view but nonetheless constitutes an interesting trade-off in a crisis management context. In fact, a post-treatment may be sufficient to greatly improve the response efficiency because the solution plans are inherently highly parallelizable as they involve many organizations. Whether one should relax assumption A7 (by taking into account the duration required to execute plan's actions) is fairly debatable. Such a feature could be interesting as it may improve the efficiency of the response (from a scheduling perspective) or ensure that the produced plans are deployable in a given time frame. Nevertheless, using temporal planning techniques is not always feasible or even desirable in a crisis management setting. Indeed, it requires to precisely estimate the execution time of each responders' actions which may be either impossible or impracticable as it would monopolize various experts in a critical step of the crisis response. Moreover, working at a strategic level, decision-makers don't produce fine-grained plans thus limiting theirs needs regarding explicit time management. One can nonetheless adopt an intermediate position by considering assumption A7 as valid while modelling the execution time of the plan as a preference. In light of this discussion, we believe that assumption A7 should not be relaxed in a crisis management context.

In a nutshell, crisis management planning problems are formalized by relaxing assumptions A1, A5, A6 and A8 of the conceptual model for planning. As produced plans may no longer be sequential, the notion of action is generalized through sets of non mutually exclusive actions denoted *happenings* (see (Fox and Long 2003) for a complete formalization). An happening A contains one or several actions that can be performed concurrently ; when executed in a state s, it leads to an other state $s' = \gamma(s, A)$. Moreover, σ_i is used to denote the real world observation corresponding to the theoretical state s_i.

Crisis Management Planning Problem. Given $\Sigma = (S, A, \gamma)$ a state-transition system, $s_0 \in S$ an initial state and $S_{G'} \subseteq S$ the set of the states verifying the extended goals G', let X denote the set of sequences of happenings $\langle A_1, \cdots, A_n \rangle$ corresponding to the sequences of states $\langle s_0, \cdots, s_n \rangle$ such that $s_1 = \gamma(s_0, A_1), \cdots, s_n = \gamma(s_{n-1}, A_n)$ and $s_n \in S_{G'}$. Let $\gtrsim$ be a partial order over X constructed from the extended preferences of the problem. A plan x is said to be an optimal solution if $\forall x' \in X, x \gtrsim x'$. The crisis management problem consist to:

1. Find a solution plan $x \in X$ given Σ, s_0 and G'. Plan x is not required to be optimal but is expected to be optimized with respect to the partial order $\gtrsim$.
2. Provide a procedure to compute a distance d_i between each state s_i of x and its corresponding observation σ_i. This procedure is used to trigger dynamic planning mechanisms whenever d_i is greater than some threshold θ.

The remainder of this study only focuses on the first part of the crisis management planning problem. The second part of the problem can be solved using approaches from the automated planning community (Ghallab et al. 2016) or from the crisis management community (Bénaben, Montarnal, et al. 2017; Barthe-Delanoë et al. 2014).

INFORMATION SYSTEM FOR PLANNING IN CRISIS MANAGEMENT CONTEXT

This section describes core features that an information system producing plans in order to support the crisis managers should possess. Such a system is intrinsically a decision support system and should not be exclusively reduced to an automated planning tool. If it were the case, the system would have little to no added value for the decision-makers. Indeed, automated planning tools process inputs expressed in a formalism unknown to the operational experts of crisis management which is named PDDL (McDermott et al. 1998) and provide no extra information regarding their outputs. These observations highlight two operational requirements that the aforementioned decision support system should fulfil. First, in order for the system to be used effortlessly, decision-makers should be able to describe the problem they are trying to solve using terms from the operational language of crisis management. Indeed, it is inconceivable to ask them to learn a scientific formalism for this purpose. Moreover, as the system aims to assist the crisis managers in their decision process, it should provide them some explanations regarding the plans it suggests. Therefore, it allows decision-makers to fully understand the recommendations of the system and decide whether they want to follow them or not.

Figure 2 describes a conceptual information system for planning in crisis management context addressing the aforementioned operational requirements. This system is used by executing three successive steps. The *modelling step* allows decision-makers to formalize the problem to solve using their operational language. Models they produce are then automatically transformed into a PDDL problem. Should they want to consider various strategies to handle the situation, the system allows decision-makers to provide several models of their preferences. During the *planning step*, a planner is used to construct a solution plan for each preference model specified by the decision-makers. In order to support the decision, each plan is evaluated against all available preference models. Finally, the *decision-making step* helps the members of the crisis cell to compare the suggested plans so they can choose the one they want to deploy.

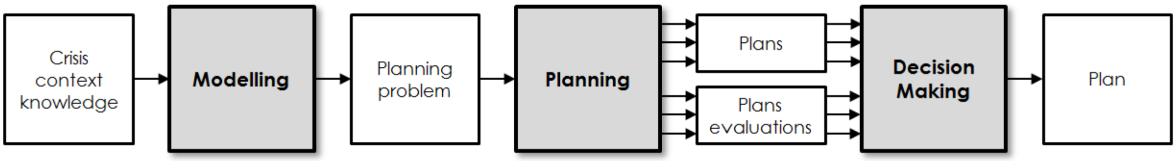

Figure 2. Aid decision system supporting decision-makers during the construction of crisis management plans

During the modelling step, decision-makers create four models (situation, organizations, goals and preferences) based on information they possess regarding the critical situation. One should note that the modelling choices made constitute a decision-making act since the solution plans are directly produced from these models. The *situation model* contains all the relevant elements of the world that should be considered thus providing a description of the initial situation of the crisis. In addition, each organization involved in the crisis response fulfils an *organization model* specifying the relevant actions it can execute as well as its available resources. The objectives that the crisis managers want to achieve are described in the *goals model* and their preferences are captured through *preferences models*. The higher the score of a plan against a preference model, the more satisfactory it is from the point of view of decision-makers. Metamodelling is the keystone of such a knowledge representation. A metamodel is a set of concepts and rules referring to those concepts used to construct models. It imposes a structure in the representation of the knowledge of the considered domain. Thus, it helps users to efficiently create models while ensuring that their models will be coherent enough to be handled by automated processing tools. Several metamodels for crisis management have been proposed, see for instance (Bénaben, Lauras, et al. 2016). Once operational models have been constructed, a *model transformation* (Truptil et al. 2010) is used in order to produce the corresponding planning problem. When a model transformation is performed, concepts from a *source metamodel* are automatically converted into concepts of a *target metamodel* using some predefined mapping rules. Such a transformation is possible since a metamodel for planning (which can be thought of as an abstraction of the PDDL language) intrinsically encompasses many concepts used in a metamodel for crisis management as explained in (Bidoux 2016).

Based on the outputs of the modelling step, the planning step produces several candidate plans as well as their respective evaluations. Decision-making step allows crisis cell's members to visualize the suggested plans. The system also provides information on the proposed plans with respect to each of the decision-makers' preferences. These information should help decision-makers to understand the plans' scores regarding the preference models considered thus allowing them to efficiently compare solution plans and take informed decisions. Moreover, crisis managers have the possibility to perform modifications on the suggested plans if necessary. Finally, if propositions made by the system are unsatisfactory, decision-makers may choose to ignore them or to modify their input models in order to ask for new ones.

Practitioners often perform the three aforementioned steps manually by writing relevant information on a flip chart. Therefore, using the proposed system should be no longer than a manual solving provided that the system's user interface is properly designed. The next section describes a proof of concept implementation of such a system.

PROOF OF CONCEPT IMPLEMENTATION

The proposed system aims to support crisis decision-makers operating at a strategic level. Designing scenarios illustrating their work that are both operationally accurate and easy to understand is a complex task. Indeed, such use cases may include hundreds of operational elements to consider, dozens of organisations to manage as well as numerous objectives to achieve and therefore can not be described concisely. In a attempt to produce the most pedagogical use case possible, we choose to illustrate the considered system at an operational level but we stress out that mechanisms detailed hereafter also remain valid at a strategic level. Hopefully, this scenario inspired from the EU-funded FP7 project SECTOR (Cinque et al. 2015) is nonetheless informative enough to highlight the main features of the proposed system.

The aforementioned scenario describes a massive flood event in Northern Europe. More precisely, it takes place in the vicinity of Roermond, a town located in the region of Limburg in the Netherlands. These events occur during a winter throughout which heavy rainfalls were recorded for several months. As a result, rivers' levels reached their highest points. Meanwhile, weather forecasts indicate that heavy precipitations are still expected. Unfortunately, Roermond area is particularly exposed in case of flooding as it is crossed by numerous lakes and rivers. The goal of this use case is to handle in the most efficient way the crisis that is going to be caused by the incoming flooding. Management of this critical situation mobilizes responders for several days. During this period, decision-makers are expected to produce several plans and use a dynamic planning approach to this end. This section only illustrates the first use of the system namely the construction of the initial rescue plan.

Decision-makers describe the crisis they have to solve by creating the *situation*, *organizations*, *goals* and *preferences models*. The system provides modelling tools that allow them to quickly represent the situation while forbidding the use of elements or relationships that are not conform with the considered metamodel. For conciseness, the creation of the situation model will be the only one illustrated hereafter (cf. figure 3). Using, for instance, the metamodel presented in (Bidoux 2016), decision-makers are permitted to use the concepts of *environment component* (blue rectangle with rounded edges), *eventuality* (purple trapezium) and *state* (green ellipse) to depict the initial situation. These concepts can be linked to one another according to the metamodel's rules. Indeed, an environment component can *generate* (gray arrow) an eventuality or *impact* (violet arrow) another environment component. In addition, environmental components are *characterized* (green arrow) by their states.

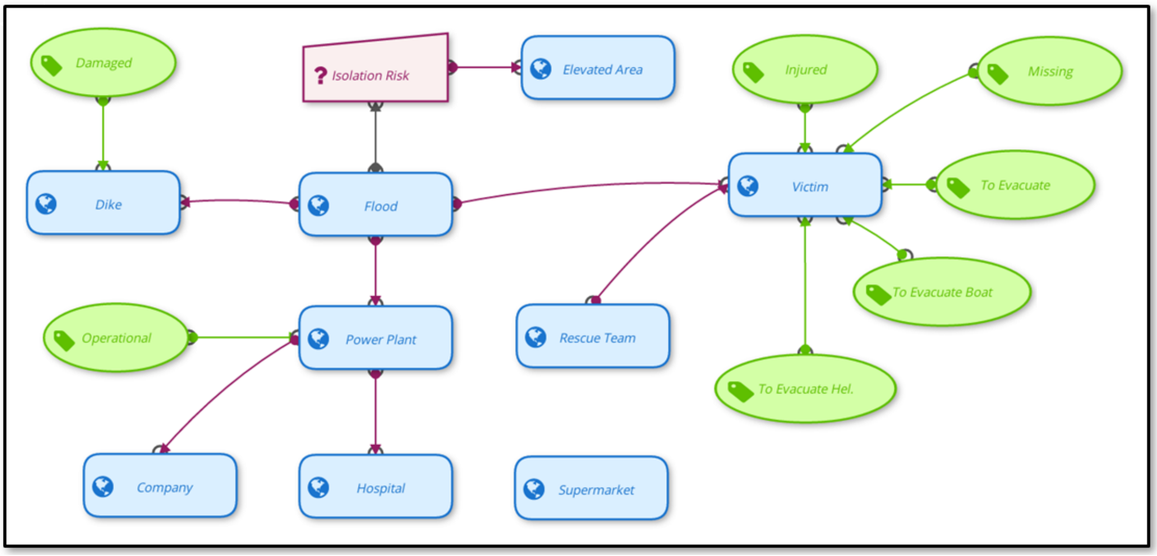

Figure 3. Situation model corresponding to the use case

Although the flooding was anticipated, it occurred sooner and faster than expected due to a climatic event whose magnitude was underestimated. As a consequence, some preventive measures have not yet been deployed at the beginning of the crisis. Unsurprisingly, the main environment component of the model is the *flood* to be handled. It damaged a *dike* and is impacting a *power plant* that currently supplies both the *hospital* of Roermond as well as production facilities of a large *company* of the region. Moreover, it also generates an *isolation risk* for the inhabitants of an *elevated area* which although non-floodable might become completely inaccessible if the water level keeps growing. In addition, the flood has engendered several *victims*, which may be *injured*, *missing* or have to be *evacuated* (by boat or helicopter in the most difficult cases). One should note that the presented situation model has been simplified for illustrative purposes as, given the modelling choices performed, one environment component should have been created for each victim.

Even if the organizations and goals models are not illustrated for conciseness, their content is briefly described hereafter. Most of involved partners are Dutch first responders (*Dutch fire department*, *Dutch police* and *Dutch Red Cross*) or local actors (*electricity provider*, *Roermond hospital*, a *local supermarket*. . .). In addition, international assistance (*German police*, *German Red Cross*, *Belgian Red Cross*) may be mobilized to strengthen Dutch forces if necessary. Each organization is asked to provide a list of the actions it can execute as well as a list of its shareable resources. In this use case, the main goal of decision-makers is to rescue the various victims of the flood. Moreover, they want to place the hospital of the town on alert and set up a temporary reception center for victims. Furthermore, the power plant must be closed as it is threatened by the flooding. Beforehand, hospital should be supplied with generator sets to ensure that it remains operational. Law enforcement operations (dissemination of safety instructions, traffic regulation. . .) are also considered in order to efficiently manage the road network of the region (priority to rescue operations and food provisioning for instance).

Preference models of the crisis managers are made with the Myriad software (Labreuche and Lehuédé 2005). The latter uses a formalism from multicriteria decision analysis in order to represent more accurately the intrinsic complexity of crisis managers' preferences. Furthermore, it allows decision-makers to construct a mathematical model of their preferences using only operational information. As a result, crisis managers are not required to have any mathematical background in order to use the Myriad software. Once the preference models are built, they can be imported into the system through a dedicated interface. In this example, decision-makers realize three preferences models that correspond to three resolution strategies they envision to implement. The first one assesses the crisis response using three criteria: (i) *efficiency* of rescue operations (modeled by the number of rescue teams deployed) ; (ii) the *comfort* of the inhabitants living in the area exposed to an isolation risk (which depends on whether food and electricity is supplied to them or not) and (iii) the *cost* of the response (excluding costs related to victim rescue). It also prohibits to request the international assistance proposed by neighboring countries so that decision-makers can assess the quality of the plan they would be able to deploy with Dutch rescue teams only. The second preference model utilizes the same criteria as the first one without forbidding the use of the international assistance. Decision-makers may therefore estimate the added value provided by the intervention of German and Belgium rescue teams. The last preference model includes an additional criterion regarding production facilities of the local company. Indeed, a brutal shutdown of the power plant could damage these production facilities thus leading to economic losses for the region in the medium term. As a result, the third model also considers a trajectory preference, denoted *company*, that reflects the will to supply the production facilities of the company using generating sets before shutting down the power plant. Production facilities could then nominally finish their production cycle before being stopped, thus avoiding damage to its machineries and infrastructures. In addition, these preference models can take into account the binary interactions that may exist between the aforementioned criteria. Two criteria may be considered as complementary (decision-makers are satisfied only if both criteria are satisfactory), substitutable (decision-makers are satisfied if at least one of the criteria is satisfactory) or independent (decision-makers satisfaction regarding one criterion is not influenced by the other one). In the first two models, *efficiency* and *comfort* are modelled as slightly complementary. In the third model, *cost* and *company* are modelled as substitutable criteria.

Once the four models describing the crisis have been realized, they can be transformed into crisis management planning problems formalized in the PDDL language thus ending the modelling step. More precisely, these problems are described using a PDDL extension which improves the language expressiveness regarding the representation of preferences. Using this extension, one can benefit from the advantages related to the aforementioned multicriteria decision analysis approach as explained in (Bidoux 2016). Throughout the planning step, a planner is used in order to solve crisis management planning problems that have been derived from crisis managers' models. Planning with preferences requires to find a good trade-off between objectives to achieve and preferences to optimize. Indeed, if solving is only focused on objectives, it is easy to find solutions but these are likely to be mediocre ones with respect to the preference model considered. On the other hand, if solving is only focused on preferences, it may even be hard to find a solution due to combinatorics of the problem. The preference-based planner ChoPlan have been used to solve the crisis management planning problems. Details regarding the heuristics it employs are outside the scope of this paper but an interested reader may refer to (Bidoux 2016). In this example, the ChoPlan planner is used to produce three solution plans as three preference models are considered by the decision-makers. Each solution plan is optimized with respect to one preference model but nonetheless evaluated against all preference models thus leading to nine evaluations in this use case. An interface permits to visualize (and modify if necessary) the proposed solutions as illustrated with plan 2 on figures 4 and 5.

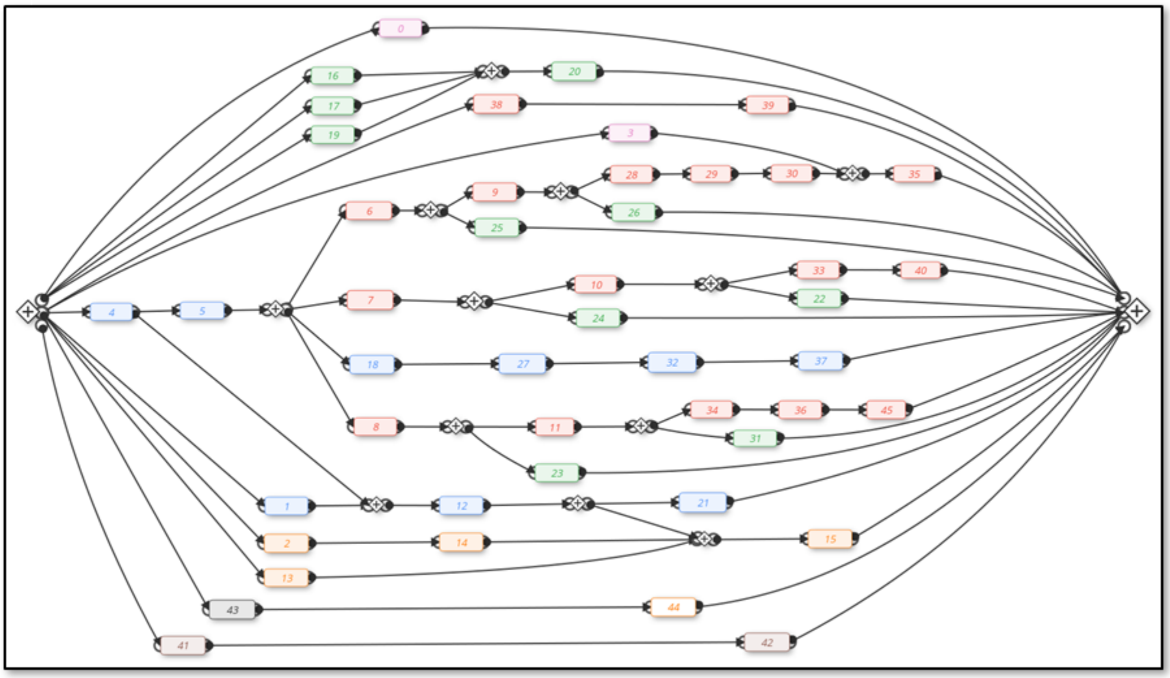

Figure 4. Plan 2 displayed in a non sequential way

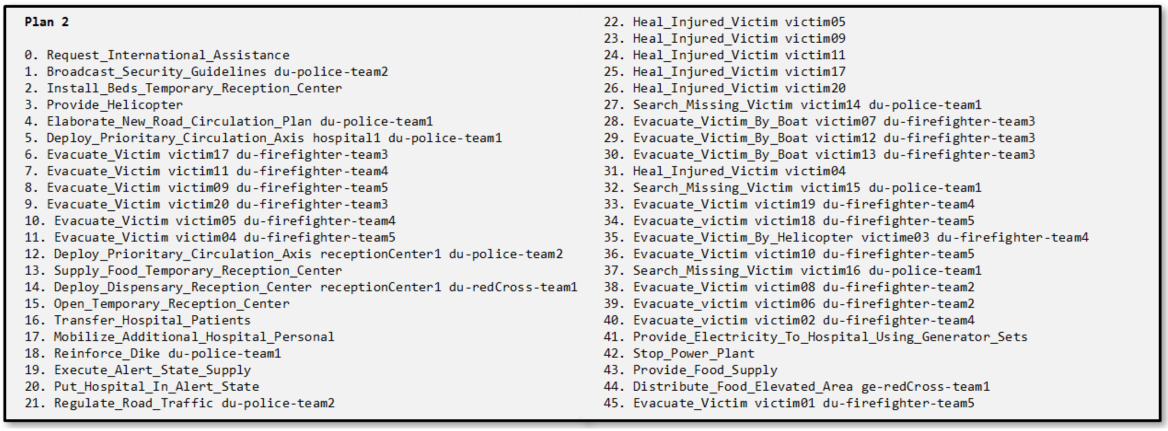

```
Plan 2

0. Request_International_Assistance
1. Broadcast_Security_Guidelines du-police-team2
2. Install_Beds_Temporary_Reception_Center
3. Provide_Helicopter
4. Elaborate_New_Road_Circulation_Plan du-police-team1
5. Deploy_Prioritary_Circulation_Axis hospital1 du-police-team1
6. Evacuate_Victim victim17 du-firefighter-team3
7. Evacuate_Victim victim11 du-firefighter-team4
8. Evacuate_Victim victim09 du-firefighter-team5
9. Evacuate_Victim victim20 du-firefighter-team3
10. Evacuate_Victim victim05 du-firefighter-team4
11. Evacuate_Victim victim04 du-firefighter-team5
12. Deploy_Prioritary_Circulation_Axis receptionCenter1 du-police-team2
13. Supply_Food_Temporary_Reception_Center
14. Deploy_Dispensary_Reception_Center receptionCenter1 du-redCross-team1
15. Open_Temporary_Reception_Center
16. Transfer_Hospital_Patients
17. Mobilize_Additional_Hospital_Personal
18. Reinforce_Dike du-police-team1
19. Execute_Alert_State_Supply
20. Put_Hospital_In_Alert_State
21. Regulate_Road_Traffic du-police-team2
22. Heal_Injured_Victim victim05
23. Heal_Injured_Victim victim09
24. Heal_Injured_Victim victim11
25. Heal_Injured_Victim victim17
26. Heal_Injured_Victim victim20
27. Search_Missing_Victim victim14 du-police-team1
28. Evacuate_Victim_By_Boat victim07 du-firefighter-team3
29. Evacuate_Victim_By_Boat victim12 du-firefighter-team3
30. Evacuate_Victim_By_Boat victim13 du-firefighter-team3
31. Heal_Injured_Victim victim04
32. Search_Missing_Victim victim15 du-police-team1
33. Evacuate_Victim victim19 du-firefighter-team4
34. Evacuate_victim victim18 du-firefighter-team5
35. Evacuate_Victim_By_Helicopter victime03 du-firefighter-team4
36. Evacuate_Victim victim10 du-firefighter-team5
37. Search_Missing_Victim victim16 du-police-team1
38. Evacuate_Victim victim08 du-firefighter-team2
39. Evacuate_victim victim06 du-firefighter-team2
40. Evacuate_victim victim02 du-firefighter-team4
41. Provide_Electricity_To_Hospital_Using_Generator_Sets
42. Stop_Power_Plant
43. Provide_Food_Supply
44. Distribute_Food_Elevated_Area ge-redCross-team1
45. Evacuate_Victim victim01 du-firefighter-team5
```

Figure 5. Plan 2 displayed in a sequential way

As part of the decision-making step, the system displays the scores of each solution plan with respect to each preference model in order to help crisis managers choosing the plan they want to deploy (see figure 6). In addition, decision-makers may complement this results overview by considering additional comparison criteria if necessary. Such criteria can be added once the planning step has been performed. They may be used to import evaluation scores regarding the proposed plans from external tools such as simulation ones for instance. Several methods have been proposed to assess action plans in emergency context, see for example (Núñez et al. 2016).

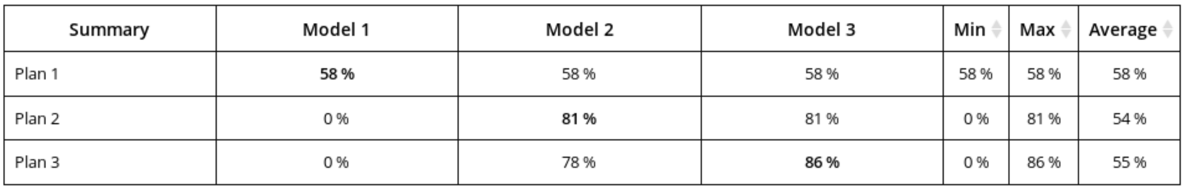

Summary	Model 1	Model 2	Model 3	Min	Max	Average
Plan 1	**58 %**	58 %	58 %	58 %	58 %	58 %
Plan 2	0 %	**81 %**	81 %	0 %	81 %	54 %
Plan 3	0 %	78 %	**86 %**	0 %	86 %	55 %

Figure 6. Scores of the solution plans in the use case

The only difference between preference models 1 and 2 concerns the constraint on the use of international assistance. The first plan (which is obtained by optimizing the first preference model) is the only solution respecting this constraint thus explaining that plans 2 and 3 have a null score with respect to preference model 1. Scores of plans 1 and 2 regarding preference model 2 can be interpreted thanks to diagrams presented in figure 7. Indeed, decision-makers have considered the response efficiency as the main criterion of this model since it is responsible for 50% of the score. In addition comfort and cost criteria both impact 20% of the score. The remaining 10% are due to a complementarity between efficiency and comfort criteria. Thus, this part of the score favors plans that are conjointly satisfactory on these two criteria. Such a case occurs for example with the second solution plan whose efficiency and comfort scores are both relatively similar and high (see bar charts in figure 7). Nevertheless, one can observe that for solution plan 1, the score of the comfort criterion is considerably weaker than the score of the efficiency criterion. The second plan differs from the first one because it resorts on the international assistance offered by Germany and Belgium. As a result, there are more rescue teams mobilized in plan 2 which impacts positively its efficiency criterion. Indeed, Red Cross teams are strengthened thus enabling this organization to carry out an additional mission in the second plan (food provisioning to the inhabitants of the area which is going to be isolated). The second plan therefore obtains a better score with respect to comfort criterion. However, its score regarding cost criterion is lower than the one of the first plan due to the costs engendered by this food distribution.

Furthermore, plans 2 and 3 can be compared with respect to preference model 3 using diagrams from figure 8. Preference model 3 presents some similarities with preference models 1 and 2 since efficiency criterion weights for 50% of the score, cost criterion represents 20% of the score and the complementarity between efficiency and comfort criteria also account for 10% of the score. Nevertheless, the remaining 20% of the score are defined by an interaction between cost and company criteria. Indeed, in this model, these two criteria are completely substitutable one to another. Thus, a good score on the company criterion can offset a bad cost of the solution and vice versa. Contrary to the second plan, the third solution (which is obtained by optimizing the third preference model) respects the preference regarding the power supply of the company's production facilities. As a result, the score of the third plan according to the company criterion is maximal. On the other hand, its score on the cost criterion is worse than the one of the second plan 2 due to the expenses engendered to supply electricity to the company. Since these two criteria are substitutable for one another, plan 3 surpasses plan 2 according to the third preference model.

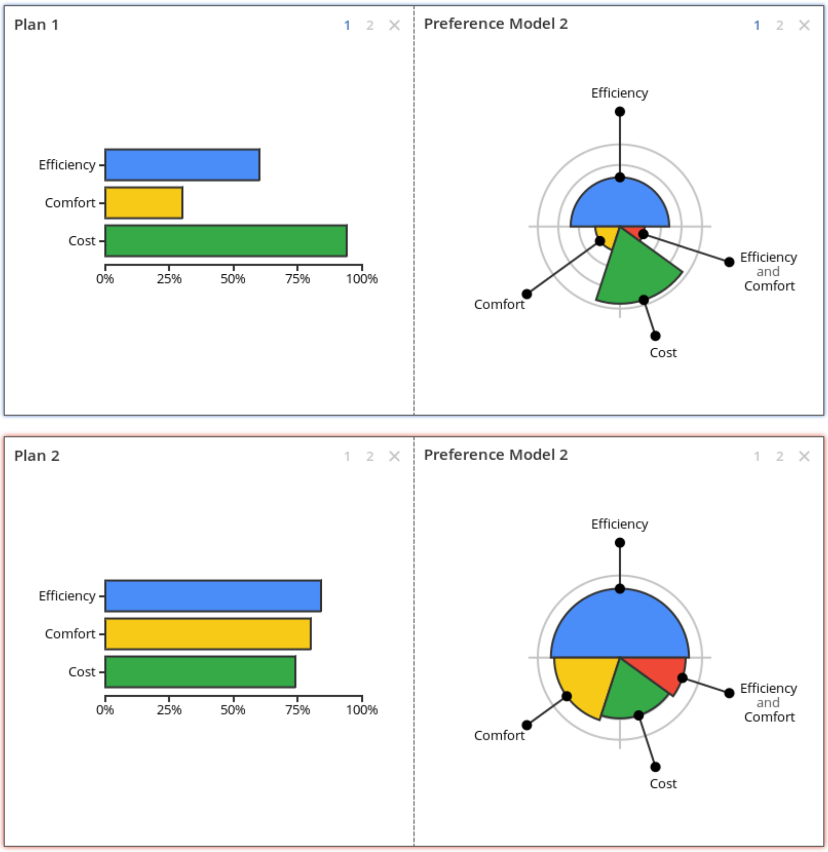

Figure 7. Comparaison of plans 1 and 2 with respect to preference model 2

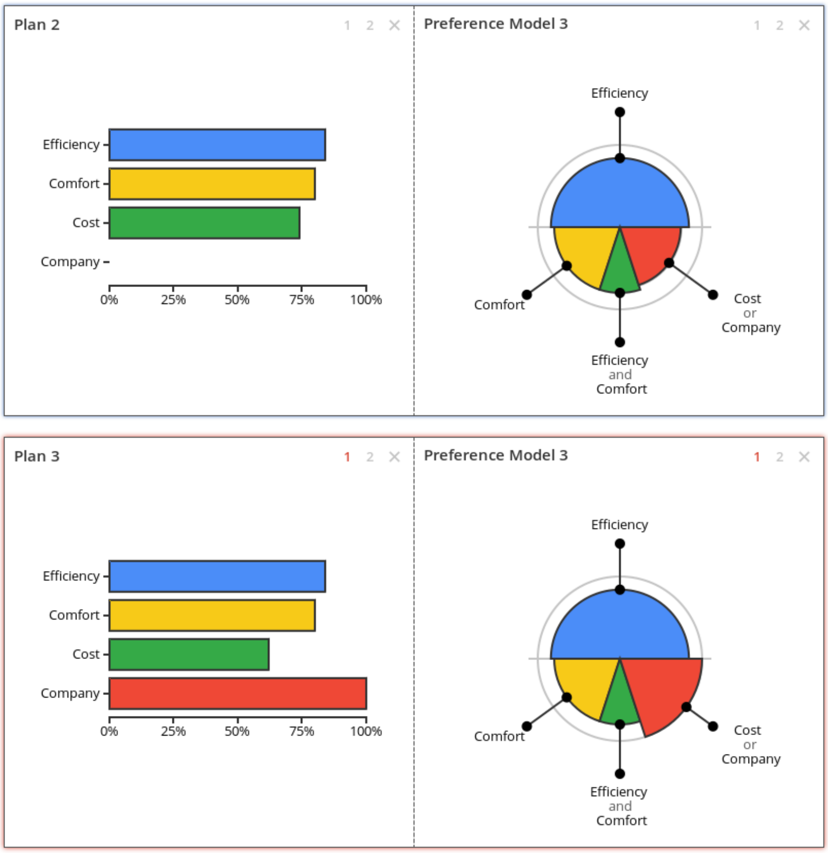

Figure 8. Comparaison of plans 2 and 3 with respect to preference model 3

In this use case, crisis managers finally choose to deploy the third plan. Indeed, it may seem unwise in this example not to resort on the aid of international partners in light of their impact on the quality of the response. Moreover, additional costs associated with the third plan seems acceptable to decision-makers given the advantages of the latter compared to the second plan.

CONCLUSION

In this paper, we have discussed the use of automated planning for crisis management. On the theoretical side, we have formalized the crisis management planning problem in light of the various assumptions of the conceptual model for automated planning. In order to address this problem, one have to (i) find a plan solving the considered problem and (ii) detect divergences between the expected situation and the real world during the plan's deployment. On the practical side, we have described some operational requirements that an operational system addressing the first part of this problem should possess. In addition, a conceptual information system satisfying these operational requirements have been proposed. A proof of concept implementation of this system has also been described through an operational use case involving a massive flood event in Northern Europe.

An interesting perspective regarding the theoretical formalization of the crisis management planning problem would be to study the case obtained when assumption A7 of the conceptual model for automated planning is also relaxed. On the practical side, a notable perspective would be to improve the modelling step of the aforementioned system. Indeed, assisting decision-makers by suggesting them elements to consider (or even by generating some parts of the models automatically) and by highlighting potential problems in their models would be a great addition. Such a feature could be implemented using inference mechanisms on an ontology populated with experts' feedbacks from past crises. To finish, the main perspective for this work would be to extend the proposed system so that it addresses the second part of the crisis management planning problem too. As a result, one may obtain a fully functional information system supporting decision-makers' planning activities in crisis management contexts.

REFERENCES

Aligne, F. and Savéant, P. (2011). "Automated Planning in Evolving Contexts: an Emergency Planning Model with Traffic Prediction and Control". In: *Proc. Future Security Conf., Bonn, Germany*, pp. 393–421.

Altay, N. and Green, W. G. (2006). "OR/MS research in disaster operations management". In: *European Journal of Operational Research* 175.1, pp. 475–493.

Barthe-Delanoë, A.-M., Truptil, S., Bénaben, F., and Pingaud, H. (2014). "Event-driven agility of interoperability during the run-time of collaborative processes". In: *Decision Support Systems* 59, pp. 171–179.

Bénaben, F., Lauras, M., Truptil, S., and Salatgé, N. (2016). "A Metamodel for Knowledge Management in Crisis Management". In: *49th Hawaii International Conference on System Sciences*, pp. 126–135.

Bénaben, F., Montarnal, A., Truptil, S., Lauras, M., Fertier, A., Salatgé, N., and Rebiere, S. (2017). "A Conceptual Framework and a Suite of Tools to Support Crisis Management". In: *Proceedings of the 50th Hawaii International Conference on System Sciences*.

Bidoux, L. (2016). "Preference-based planning using the Multi-Attribute Utility Theory along with a Choquet integral: application to organizations' interoperability in crisis management". PhD thesis. Ecole Nationale Supérieure des Mines d'Albi - Université de Toulouse.

CCA (2014). "Business continuity structured glossary". In: *Livre blanc du Club de la Continuité d'Activité*.

Choquet, G. (1953). "Theory of capacities". In: *Annales de l'Institut Fourrier* 5, pp. 131–295.

Cinque, M., Esposito, C., Fiorentino, M., Mauthner, J., Szklarski, Ł., Wilson, F., Semet, Y., and Pignon, J.-P. (2015). "SECTOR: Secure Common Information Space for the Interoperability of First Responders". In: *Procedia Computer Science* 64, pp. 750–757.

Dean, T. and Wellman, M. (1991). *Planning and Control*. Morgan Kaufmann Publishers.

Dyer, J. S. (2005). "MAUT - Multiattribute Utility Theory". In: *Multiple Criteria Decision Analysis: State of the art surveys*. Springer, pp. 265–292.

Eisenhower, D. D. (1957). *Remarks at the National Defense Executive Reserve Conference, 14 Novembre 1957. The American Presidency Project*. `http://www.presidency.ucsb.edu/ws/?pid=10951`.

Fox, M. and Long, D. (2003). "PDDL2.1: An Extension to PDDL for Expressing Temporal Planning Domains". In: *Journal of Artificial Intelligence Research* 20, pp. 61–124.

Gerevini, A. and Long, D. (2006). "Preferences and Soft Constraints in PDDL3". In: *ICAPS workshop on planning with preferences and soft constraints*, pp. 46–53.

Ghallab, M., Nau, D., and Traverso, P. (2004). *Automated Planning: Theory and Practice*. Elsevier.

Ghallab, M., Nau, D., and Traverso, P. (2016). *Automated Planning and Acting*. Cambridge University Press.

Glasser, R. and Guha-Sapir, D. (2016). *Poverty and Death: Disaster Mortality 1996-2015*.

Grabisch, M. and Labreuche, C. (2010). "A decade of application of the Choquet and Sugeno integrals in multi-criteria decision aid". In: *Annals of Operations Research* 175.1, pp. 247–286.

Labreuche, C. and Lehuédé, F. (2005). "MYRIAD: a tool suite for MCDA". In: *EUSFLAFT*. Vol. 5, pp. 204–209.

McDermott, D., Ghallab, M., Howe, A., Knoblock, C., Ram, A., Veloso, M., Weld, D., and Wilkins, D. (1998). *PDDL - The Planning Domain Definition Language*. Tech. rep. CVC TR-98-003/DCS TR-1165.

Núñez, A.-G., Penadés, M. C., and Canós, J. H. (2016). "QuEP: Building a Continuous Improvement of Emergency Plans Management". In: *13th International Conference on Information Systems for Crisis Response and Management (ISCRAM 2016)*.

Treurniet, W., Van Buul-Besseling, K., and Wolbers, J. (2012). "Collaboration Awareness. A necessity in crisis response coordination". In: 9^{th} *International Conference on Information Systems for Crisis Response and Management (ISCRAM 2012)*.

Truptil, S., Bénaben, F., Salatgé, N., Hanachi, C., Chapurlat, V., Pignon, J.-P., and Pingaud, H. (2010). "Mediation Information System Engineering for Interoperability Support in Crisis Management". In: *Enterprise Interoperability IV*. Springer, pp. 187–197.

Van De Walle, B. and Turoff, M. (2007). "Emergency Response Information Systems: Emerging Trends and Technologies". In: *Communications of the ACM* 50.3, pp. 29–31.