Operational Emergency Response under Informational Uncertainty: A Fuzzy Optimization Model for Scheduling and Allocating Rescue Units

Felix Wex
Albert-Ludwigs-Universität Freiburg
felix.wex@is.uni-freiburg.de

Guido Schryen
Universität Regensburg
guido.schryen@wiwi.uni-regensburg.de

Dirk Neumann
Albert-Ludwigs-Universität Freiburg
dirk.neumann@is.uni-freiburg.de

ABSTRACT

Coordination deficiencies have been identified after the March 2011 earthquakes in Japan in terms of scheduling and allocation of resources, with time pressure, resource shortages, and especially informational uncertainty being main challenges. We suggest a decision support model that accounts for these challenges by drawing on fuzzy set theory and fuzzy optimization. Based on requirements from practice and the findings of our literature review, the decision model considers the following premises: incidents and rescue units are spatially distributed, rescue units possess specific capabilities, processing is non-preemptive, and informational uncertainty through linguistic assessments is predominant when on-site units vaguely report about incidents and their attributes, or system reports are not exact. We also suggest a Monte Carlo-based heuristic solution procedure and conduct a computational evaluation of different scenarios. We benchmark the results of our heuristic with results yielded through applying a greedy approach. The results indicate that using our Monte Carlo simulation to solve the decision support model inspired by fuzzy set theory can substantially reduce the overall harm.

Keywords

Decision Support Systems, (Fuzzy) Optimization, Coordination, Informational Uncertainty, Fuzzy Set Theory.

INTRODUCTION

Natural disasters, including earthquakes, tsunamis, floods, hurricanes, and volcanic eruptions, have caused tremendous harm and continue to threaten millions of humans and various infrastructure capabilities each year. Being consistent with the terminology of the International Federation of Red Cross and Red Crescent Societies (IFRC) and the U.S. Federal Emergency Management Agency (FEMA), we use the term "disaster" in the following sense (IFRC): "A disaster is a sudden, calamitous event that seriously disrupts the functioning of a community or society and causes human, material, and economic or environmental losses that exceed the community's or society's ability to cope using its own resources." In this study, we focus on disasters based on natural disasters, not on technological, man-made, or attack-based disasters. In contrast to the latter types, their natural counterparts are not preventable. Thus, the actions that need to be taken before, during, and after disasters and the used data are different. For example, risk management of floods and hurricanes can draw on geological data, while the risk management of nuclear attacks by terrorists cannot do so.

The coordination of resources during natural disasters is characterized by a high level of informational uncertainty due to the chaotic situation, severe resource shortages, and a high demand for timely information in the presence of the disruption of infrastructure support (Chen, Sharman, Rao, and Upadhyaya, 2008). The March 2011 earthquakes near the coast of Sendai, Japan manifested these presumptions, as did the management of the succeeding nuclear disaster (Krolicki, 2011). Emergency operations centers (EOC) were confronted to the partial breakdown of information systems and uncertain access infrastructures to incidents (roadblocks). Officials had to deal with numerous incidents where more than 27,000 people were found dead or missing and some 150,000 Japanese displaced (Sanders, 2011). An improvisation and decentralization of the actions of local commanders and rescue teams have been noticed. The involvement of numerous, international organizations with different disaster response policies, resources, and technological infrastructures and capabilities led to distributed planning and implementing of response actions (Chawla, 2011). Poor communication between geographically dispersed EOCs, a lack of accurate data, and an immense time pressure intensified the dilemma

(Deutsche Presse-Agentur, 2011; Dmitracova, 2010). Even though resource scarcity can occur, we argue that the "appropriate allocation of [spatially distributed] resources is more important (…) [and] a problem of coordination" (Comfort, Ko, and Zagorecki, 2004; Klingner, 2011).

The above issues reveal that the automation of allocating rescue units to incidents remains a challenge in utilizing Emergency Response Systems (ERS). In practice, as told by associates of the German Federal Agency of Technical Relief (THW), assignments and schedules for resources are still derived through the application of greedy policies: for example, based on a ranking of incidents in terms of destructiveness, the most severe incidents are sequentially handled by the closest, idle rescue units, (also stated by Comfort, 1999). However, this rather straightforward – albeit in many cases common and favorable – rule ignores estimated processing times of incidents, which are often not exact but may significantly affect casualties of less severe incidents and thereby the resulting harm.

When EOCs face the challenge to coordinate their rescue units, they usually find a chaotic situation in which much information is uncertain. For example, the severity of incidents is described in terms of linguistic terms, such as "lots of damage" or "a little fire burning". Subsequently, information on how much time rescue units need to process these incidents is vague, too. The chaotic situation does also not allow making precise statements on how long rescue units travel between two points of incidents as the traffic infrastructure may be severely affected. All these types of information have in common that the impreciseness of predictions is due to a lack of information, belief, and linguistic characterizations, which all are deemed some of the most important roots of uncertainty (Zimmermann, 2000). In the absence of statistical information and in the presence of subjective uncertainty we account for these roots of uncertainty by drawing on fuzzy set theory (Zadeh, 1965) among the many available uncertainty theories. Fuzzy set theory in emergency response situations has been stated appropriate (Altay and Green III, 2006). This fact is based on the idea that "a [fuzzy set theory based] framework provides a natural way of dealing with problems in which the source of imprecision is the absence of sharply defined criteria of class membership rather than the presence of random variables." (Zadeh, 1965)

We also argue that time is the most crucial factor during emergency response coordination and argue for the primary goal to minimize completion times of incidents, where completion times can be defined as the duration of the occurrence of an incident until its extinction. As the literature provides only few papers on decision support in emergency response situations that deal with the assignment of the rescue units to incidents by using optimization modeling (see next section), the purpose of our paper is to suggest a novel mathematical decision model, and to propose and to evaluate a solution heuristic (a Monte Carlo simulation).

The remainder of this paper is structured as follows: based on a review from scholars and interviews with practitioners, we identify requirements for a decision support model in the subsequent section. Thirdly, we present our artifact, a quantitative decision model incorporating findings from Fuzzy (Set) Theory. Then we describe the computational evaluation, which attests the advantages of the suggested solution approach over a procedure which is found in practice. The paper closes with a conclusion.

REQUIREMENT ENGINEERING

Our requirement engineering section considers two sources: first, in order to account for the experience of practitioners, we interviewed associates of the German Federal Agency for Technical Relief (THW), who were in direct contact with the first German search and rescue teams after the major earthquakes in Japan in March 2011 and who were knowledgeable with respect to on-site coordination. Second, we use knowledge and experience of scholars with domain expertise (literature review).

The issue of allocating and scheduling rescue units during emergency response has been addressed only rarely in the literature. (Fiedrich, Gehbauer, and Rickers, 2000), (Rolland, Patterson, Ward, and Dodin, 2010), and (Wex, Schryen, and Neumann, 2011) all attest that rescue units' assignments and schedules are an understudied, yet highly relevant topic for IS research, and they suggest applying decision optimization models in a centralized manner, with a particular focus on the allocation of distributed rescue units to incidents. However, (Rolland et al., 2010) neglect the fact that rescue units are diverse in their skills. (Fiedrich et al., 2000) consider only one type of incident: earthquakes. (Wex et al., 2011) take heterogeneous rescue units into account for coordination in a centralized way and they do not concentrate on one distinct disaster type only. Yet, what all three aforementioned works lack is that they propose no possibilities to handle informational uncertainty, i.e. due to linguistic assessments, even though all authors identify this issue as one crucial challenge. In the autonomous agents community, several works have been proposed that handle task allocation in uncertain environments mainly by using auctions that either do not explicitly coordinate rescue agents or that do not fully consider the characteristics of the emergency response domain (Nair, Ito, Tambe, and Marsella, 2002; Ramchurn, Rogers, Macarthur, Farinelli, Vytelingum, Vetsikas and Jennings, 2008).

The first 72 hours after any catastrophe, the so-called critical deadline, are essential for surviving (Engelmann and Fiedrich, 2007; Reijers, Jansen-Vullers, Zur Muehlen, and Appl, 2007). Therefore, any research presenting quantitative artifacts must demonstrate its ability to (re-)act timely in real-world applications. As a consequence, any decision support system has to provide allocation and scheduling suggestions that are not only practically feasible and justifiable (in terms of specific criteria to be defined) but that are also made speedily available to aid organizations. Thus, we define Requirement 1:

Timeliness of decision support

A lack of centralized coordination may yield deficiencies in terms of control over actions of units and error-prone supervision caused by inhomogeneous or duplicate commands to multi-autonomous agents with limited information about other actors' status and positions (Airy, Mullen, and Yen, 2009). When international aid organizations come and work together during a disaster, they consequently "put themselves under the control of the responsible EOC without losing their internal, autarkic command structure" (cit. THW, translated). Following the argument of (Rolland et al., 2010), that congruent activities and non-interference among multiple decision-makers are ensured by separating operational areas, we further argue that by installing a decision support system for single, closed operational areas or jurisdictions, computer assistance is more consistent, penetrative and thus more effective. This is particularly important for situations when single organizations "are assigned their own operational area, which is then to be operated independently such that the organization acts as an EOC" (cit. THW, translated). Since key planning tasks in disaster management also include the assignment of distributed rescue units to incidents and their scheduling, it is essential to consider information on capabilities, processing times, and travel times of rescue units as well as information on incident locations and types, albeit information may be uncertain and not crisp (Fiedrich et al., 2000). We need to assume that commanders do have such (complete) information on their own resources and (at least uncertain) information on incidents to process. This argument is reasonable when (geographic) information systems are optimally installed and interconnected, suitable protocol standards exist as well as communication between commanders and on-site agents is distinct and stable. To sum up, any decision support system for resource assignments and scheduling needs to have such complete albeit uncertain information on resources and incidents available, and the commander in charge that applies the decision support system can decide and act autonomously. Requirement 2 identifies:

Completeness of (decentralized) information and decision autonomy of decentralized commanders

"When several, differently-skilled rescue teams collaborate, it is often hard to strictly classify their structure, capabilities, and their behavior. In fact, rescue units are diverse in their capabilities and sizes. (…) Generally, incidents are classified into types, such that distinct specialized rescue units are required, although it is more than challenging to prioritize a scene and to tell when search-and-rescue or firefighting brigades need to be demanded." (cit. THW, translated) Accounting for this insight of practitioners, we argue that decision support systems need to consider heterogeneous types of incidents and distinct capabilities of rescue units. For example, units can be paramedics, fire brigades, or policemen. In cases where no detailed information is available, it seems straightforward to classify incidents coarse-grained and to assign one of the rescue units that is deemed most appropriate for addressing the incident. In other cases, more detailed information on incidents is available and can be matched with specific capabilities of rescue units. Requirement 3 reads:

Consideration of differently skilled rescue units and heterogeneous incidents

As mentioned above, during any natural disaster much information remains uncertain: "decision support systems used in disaster management must cope with the complexity and uncertainty involved with the scheduling assignment of differentially-skilled personnel and assets to specific tasks" (Rolland et al., 2010). Thus, commanders often face uncertain, unconfirmed, and even contradictory information (Comes, Conrado, Hiete, Kamermans, Pavlin, and Wijngaards, 2010). This also includes statements on the severity levels of incidents. As characteristics of incidents are often described and assessed by humans, linguistic estimations are common. Other issues of (linguistic) uncertainty include non-accurate approximations about processing times, distances between incidents and positions of rescue units. Thus, we argue that decision support systems need to account for linguistic, non-probabilistic informational uncertainty, Requirement 4:

Consideration of linguistic informational uncertainty

We recall that uncertainty in chaotic emergency situations occurs due to incomplete and imprecisely stated information and not due to statistical uncertainty. Consequently, we do not suggest a probabilistic optimization model but a decision model that draws on fuzzy set theory, fuzzy arithmetic, and fuzzy optimization. The following description of the decision model briefly introduces into the key concepts; for a comprehensive overview of these areas, see (Buckley and Eslami, 2002; Klir and Yuan, 1995).

A FUZZY DECISION MODEL

Fuzzy Set Theory

Fuzzy set theory generalizes traditional set theory by providing for a degree of membership that indicates if an element belongs to a fuzzy set, in contrast to (crisp) set theory, wherein an element explicitly either comes with a set or not. A specific type of a fuzzy set is a fuzzy number (Buckley and Eslami, 2002), which is formally defined by $\{(x, \mu_{\tilde{N}}(x)) | x \in R\}, \mu_{\tilde{N}}: R \rightarrow [0,1]$, where $\tilde{N}$ is referred to as fuzzy number. $\mu_{\tilde{N}}$ is denoted as the membership function of $\tilde{N}$, and it outputs the degree with which $x \in R$ belongs to $\tilde{N}$. For example, the fuzzy number $\widetilde{10}$ which is to be equivalently seen as "real numbers close to ten" may be given by the membership function $\mu_{\widetilde{10}}(x) = (1 + (x - 10)^{-2})^{-1}$ ($x \in R^{\neq 10}$), $\mu_{\widetilde{10}}(10) = 1$. Note that the membership function differs from a probability density function in two regards: $\int_{-\infty}^{\infty} \mu_{\tilde{N}}(x)\, dx$ does not need to equal 1, and it mirrors the subjective attitude of an individual rather than reflecting statistical evidence. This is advantageous in cases where probabilities or exact data is not available, but subjective estimates of experienced experts are given. In the emergency response setting such cases are typically prevalent. The Fuzzy Decision Model makes use of the concept of symmetric triangular fuzzy numbers. A triangular fuzzy number N=(a,b,c) , a<b<c, {a,b,c}∈R, is a fuzzy set over R, with the membership function

$$\mu_N(x) = \begin{cases} \frac{x-a}{b-a}, if\ a \le x < b \\ \frac{c-x}{c-b}, if\ b \le x \le c \\ 0, otherwise \end{cases}$$

If *l*:=(c-b)=(b-a), the triangular fuzzy number is symmetric. We use symmetric fuzzy numbers with *l*=0.1*b depending on the degree of uncertainty we are facing. This corresponds to "10% fuzziness". While for many crisp optimization problems algorithms are available, this is not true for fuzzy optimization problems (Buckley and Jowers, 2008). Thus, we apply a Monte Carlo simulation for the computational evaluation in the follow-up.

Problem Description

The model is designed to schedule and assign various rescue units to incidents. It favors commanders with decision autonomy by delivering allocation solutions and schedules for all rescue units employed. The evolving question is how these units can be scheduled and assigned to incidents such that the sum of all completion times, which are individually multiplied by the individual factors of destruction, can be minimized. Factors of destruction indicate the (ordinal) levels of severity of incidents.

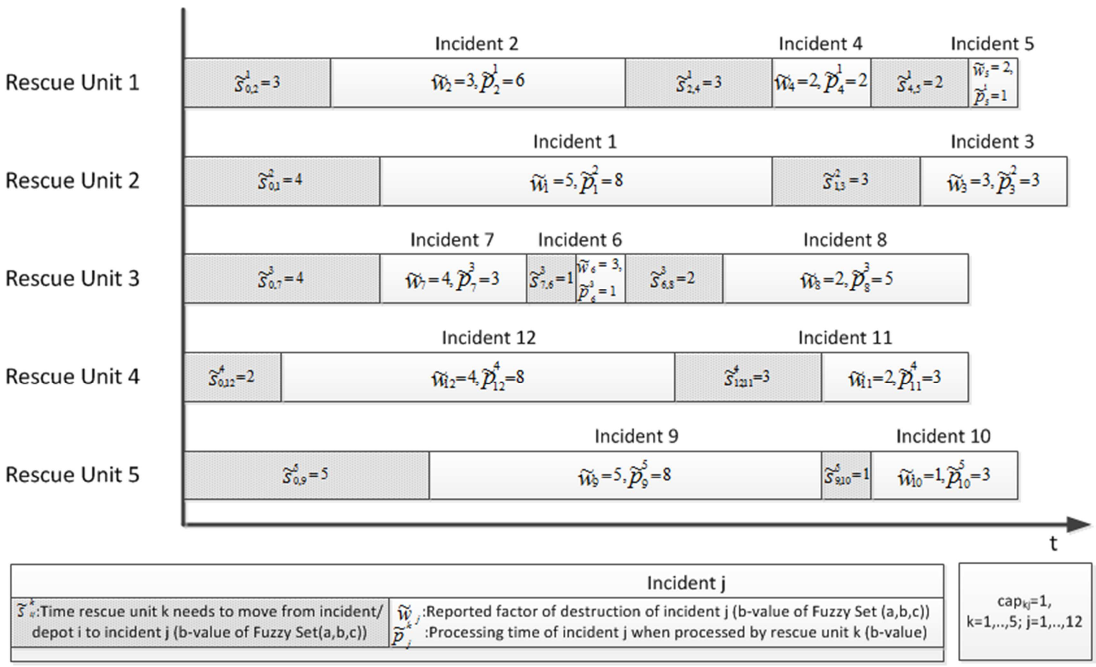

Figure 1. Desideratum: optimal schedules and assignments

We consider a situation in which the number of available rescue units is lower than the number of incidents that need to be processed. This ratio accounts for a typical natural disaster situation: *"During any large-scale disaster, there tend to be more incidents than rescue units. This is especially true within those critical minutes of* *the chaos phase." (cit. THW, translated)* An incident can be processed by a rescue unit only if this rescue unit features the specific capability that is required to process this incident. Two types of time spans are relevant: a) travel times that rescue units need to travel between two incident locations, b) processing times. We illustrate the problem description in Figure 1, which shows a feasible and valid solution of a problem instance with 5 rescue units and 12 incidents. In this instance, the level of severity (factor of destruction w_j) of incidents varies between 1 and 5. The sample schedule accounts for the specific requirements of the incidents as each rescue unit k features the respective capability that is required by incident j (cap_{kj}=1).

Mathematical Model

We define completion times as the sum of the processing times and the time the incident had to "wait" until being processed by a qualified rescue unit. This "waiting time" consists not only of processing times of incidents that have been processed previously by the assigned unit but also of the time needed to move from one incident to be processed to the next. In the assumed setting, we propose that a) the multiplication of completion times and factors of destruction $\widetilde{w_j}$ is an adequate proxy for the quality of emergency response, b) each incident can be processed by at most one unit at a time with each unit processing at most one incident at a time, c) processing is non-preemptive, and d) some data (processing times $\widetilde{p_i^k}$, severity of incidents $\widetilde{w_j}$, and travel times $\widetilde{s_{ij}^k}$) is available, deterministic, but highly uncertain and therefore not crisp. A discussion of these assumptions is included in our conclusions. Summarizing the restrictions and requirements from above, this decision model can be formulated as a non-linear binary optimization model. The mathematical formulation is provided below:

$$\min \sum_{j=1}^{n} \widetilde{w_j} \left(\sum_{i=0}^{n} \sum_{k=1}^{m} \left[\widetilde{p_i^k}\, Y_{ij}^k + \left(\widetilde{p_j^k} + \widetilde{s_{ij}^k} \right) X_{ij}^k + Y_{ij}^k \left(\sum_{l=0}^{n} X_{li}^k \widetilde{s_{li}^k} \right) \right] \right) \quad \text{(O)}$$

s.t.

$$\sum_{i=0}^{n} \sum_{k=1}^{m} X_{ij}^k = 1 \,, \quad j = 1,..,n \quad \text{(C1)}$$

$$\sum_{j=1}^{n+1} \sum_{k=1}^{m} X_{ij}^k = 1 \,, \quad i = 1,..,n \quad \text{(C2)}$$

$$\sum_{j=1}^{n+1} X_{0j}^k = 1 \,, \quad k = 1,..,m \quad \text{(C3)}$$

$$\sum_{i=0}^{n} X_{i(n+1)}^k = 1 \,, \quad k = 1,..,m \quad \text{(C4)}$$

$$Y_{il}^k + Y_{lj}^k - 1 \le Y_{ij}^k \,, \quad i = 0,..,n\,;\ j = 1,..,n+1; k = 1,..,m; l = 1,..,n \quad \text{(C5)}$$

$$\sum_{i=0}^{n} X_{il}^k = \sum_{j=1}^{n+1} X_{lj}^k \,, \quad l = 1,..,n\,; k = 1,...,m \quad \text{(C6)}$$

$$X_{ij}^k \le Y_{ij}^k \,, \quad i = 0,..,n; j = 1,..,n+1; k = 1,..,m \quad \text{(C7)}$$

$$Y_{ii}^k = 0, \quad i = 0,...,n+1; k = 1,..,m \quad \text{(C8)}$$

$$\sum_{j=1}^{n+1} X_{ij}^k \le cap_{kj} \,, \quad i = 1,..,n;\ k = 1,..,m \quad \text{(C9)}$$

$$X_{ij}^k, Y_{ij}^k \in \{0,1\} \,, \quad i = 0,..,n; j = 1,..,n+1; k = 1,..,m \quad \text{(C10)}$$

$$cap_{kj} \in \{0,1\}, \quad k = 1,..,m;\ j = 1,..,n \quad \text{(C11)}$$

$$\widetilde{w_j}, \widetilde{p_j^k}, \widetilde{s_{ij}^k} \in \widetilde{R^{\ge 0}} \quad \text{(C12)}$$

In addition to the real incidents 1,..,n we need to add two fictitious incidents '0' and 'n+1' with $\widetilde{p_0^k} = \widetilde{p_{n+1}^k} = 0$, and $\widetilde{s_{0j}^k}$ to be the estimated time that rescue unit k needs to move from its starting location (defined as incident i=0) to the location of incident j, and $\widetilde{s_{j(n+1)}^k} = 0$ for all rescue units k. The objective function (O) of the model minimizes the total weighted completion times over all incidents. Two decision variables X_{ij}^k and Y_{ij}^k are introduced indicating a mediate or immediate predecessor relationship between i and j when processed by rescue unit k. $\widetilde{w_j}$ is the reported factor of destruction of incident j and is modeled as a triangular fuzzy number.

Consequently, the lower the factor of destruction, the less severe is the incident. An explanation of the other mathematical terms used is provided in Table 1.

Decision Variable	***Interpretation***
X_{ij}^{k}	$X_{ij}^{k} = 1$ if incident i is processed immediately before incident j by rescue unit k, and 0 otherwise
Y_{ij}^{k}	$Y_{ij}^{k} = 1$ if incident i is processed before incident j by rescue unit k, 0 otherwise
Fuzzy Parameters	***Interpretation***
$\widetilde{p_i^k}$	Processing time that unit k needs to process incident i, $\widetilde{p_i^k} = \infty$ if rescue unit k is not capable of processing incident i
$\widetilde{s_{ij}^k}$	Travel time that unit k needs to move from location of incident i to location of incident j
$\widetilde{w_j}$	Reported factor of destruction of incident j equivalent to the severity level of an incident
Crisp Parameter	***Interpretation***
cap_{kj}	$cap_{kj} = 1$ if rescue unit k is capable of addressing incident i, and 0 otherwise

Table 1. Explanation of mathematical terms

Constraint (C1) ensures that for each real incident there is exactly one incident that is processed immediately before. Similarly, (C2) ensures that for each real incident there is exactly one incident that is processed immediately thereafter. Constraints (C3)-(C4) guarantee that in a feasible solution each rescue agent starts processing the fictitious incident 0 and ends processing the fictitious incident n+1, respectively. (C5) accounts for the transitivity criterion of any predecessor relationship. Yet, if an immediate predecessor for a specific incident 'l' exists, there also has to be a successor (C6). (C7) indicates that an immediate predecessor is a general predecessor. (C8) prohibits a reflexive, direct or indirect predecessor relationship. (C9) ensures that a rescue unit that is assigned to an incident possesses the required, incident-specific capability. (C10) makes the model a binary program. (C11) declares if a rescue unit is capable of operating an incident or not. (C12) defines all other parameters used. Each feasible solution of the minimization model represents a valid schedule and assignment for all units.

This Fuzzy Decision Model is especially able to manage informational overload and linguistic uncertainty by integrating fuzzy parameters (Requirement 4): impreciseness in reports from on-site forces is prevalent when determining travel and processing times, as well as severity of incidents. Furthermore, the model is also apt to assist (decentralized) commanders with decision autonomy but does not require exact (crisp) but complete information about all parameters (Requirement 2). In the adjacent sections, it will be shown that the model is adequate to deliver timely results within decent runtimes when applying the solution heuristic (Requirement 1).

The idea to search for something optimal during any disaster is questionable and can be doubted, especially when integrating uncertain information (fuzzy parameters) into the model. We therefore talk about the quest for the most effective allocations of rescue units in an uncertain setting. Disaster situations are evolving very fast sometimes (based on incoming information about the situation, incoming new resources, or on status changes of existing resources). Even though the presented approach seems to not account for this inherent dynamic and to be static, we *explicitly* suggest running the optimization of weighted completion times anew once other incidents appear or rescue units become idle (continuous optimization process). This way, alternatives and decisions can also be revisited and alterations be integrated.

COMPUTATIONAL EVALUATION

The Fuzzy Decision Model is a generalization of the machine scheduling problem "Identical parallel machine non-preemptive scheduling with minimization of sum of completion times" (Blazewicz, Dror, and Weglarz, 1991). Since any solution of the machine scheduling problem is computationally inefficient and thus NP-hard, so is its transformation to the Fuzzy Decision Model (see Appendix). As we face instances in practice, that need to be tackled as fast as possible, we suggest a Monte Carlo simulation as heuristic method. In the absence of knowledge of optimal solutions, we do not know lower bounds for the minimization instances, but we know solutions that would result from applying a greedy heuristic. Recapitulating the greedy approach, we assume that the most severe incident is assigned to the closest, idle rescue unit. The evaluation of all Monte Carlo results is based on the comparison with this benchmark indicating the proportionate reduction of harm. Implementations were written in the numerical computing environment MATLAB.

Data Generation and Origin

Information for conducting our simulation runs and for configuring data settings was integrated into scenarios in accordance to the THW interviews. We consider short travels as for most incidents occur in overcrowded areas, such as cities, where rescue units are placed relatively close to incidents. We also assume that processing times exceed travels, due to the hypothesis that urban areas are endangered more often than rural areas which result in lower ratios of travel to processing times. The factors of destruction indicate levels of severity and express five different stages for each incident, we refer to the advisory system concerning threat conditions and risks introduced by the U.S. Department of Homeland Security: low (1), guarded (2), elevated (3), high (4), and severe (5) harm. We tackle informational uncertainty by assigning vague linguistic assessments of any incidents to one of the ordinal factors of destruction expressed as fuzzy numbers. This assignment takes place when linguistic reports about incidents, casualties, and damage need to be processed and incidents need to be prioritized. For example, reports about incidents indicating "little damage" or "minor wounds" are assigned a fuzzy number close to 1 (low). Whereas reports about "injured, trapped, or exposed people" will be classified by a much higher fuzzy number, i.e. 5 (severe). Please note that any assignment of terms to a fuzzy number is never distinct (and thus fuzzy) and is always depending on the specific situation of the disaster. We consider this prioritization of incidents to be crucial.

The number of capabilities of rescue units is set to five, due to the difficulty to clearly classify all skilled units during any disaster (i.e. paramedics, fire brigades, police enforcement, military forces, or volunteers with various other skills). We extend the classification of rescue units stated by (New South Wales Government). Any rescue unit possesses one distinct capability each. The fulfillment of Requirement 3 can be guaranteed by this random assignment of rescue units' capabilities mapping real world scenarios with heterogeneous incidents that require specialized forces.

In all simulation runs executed, we explicitly suppose an ad-hoc, iterative operation mode and thus a continuous optimization, as noted in the above section. Since we motivated our decision model to work for (decentralized) commanders with decision autonomy within their own operational area (Requirement 2), we assume the number of incidents (within this area) to not exceed 200 at a time. We assume the number of rescue units, which are available and idle at the time of optimization, to not exceed 20. Since we advise to keep planning intervals short, we may execute an iterative Monte Carlo simulation. An iterative Monte Carlo simulation (continuous optimization respectively) is equivalent to the repeated dispatch of assignments and schedules for rescue units. Therefore, our approach is apt to react fast and timely (Requirement 1) to any upcoming situational change during the disaster. For our computational evaluation we regard seven different scenarios that result from permuting numbers of incidents and/or units.

Parameter	***Value, Range, Distribution***	***Rationale***
Rescue units	{10,20}	Realistic numbers of rescue units and incidents within operational areas
Incidents	{20,50,100, 200}	
Processing times $\widetilde{p_{k\iota}}$	Normally distributed: μ=20, σ=10	Occurrence of disasters close to overcrowded areas (thus: low travel times between incidents); WLOG: significant endurance of (mean) processing times to (mean) travel times (factor: 20:1)
Travel times $\widetilde{s_{k\iota}}$	Normally distributed: μ=1, σ=0.3	
Factors of destruction $\widetilde{w_\iota}$	Random Integer: {1,..,5}	Distinct risk levels introduced by the U.S. Department of Homeland Security
Capabilities $\{A_1, .., A_n\}$ n=5	A1: Search and Rescue A2: Paramedics / Medical Retrieval A3: Fire Brigades A4: Police Units A5: Special Casualty Access Team	Distinction of units' types and skills extending the classification of (New South Wales Government)
Iterations	1000	No significant improvements in the objective value beyond this point

Table 2. Settings in randomly generated scenarios

Results

We benchmarked all Monte Carlo simulation results to the results generated by a greedy policy (however, we did not determine optimal solution values or lower bounds): the most severe incident is assigned the closest, idle rescue unit and the remaining idle rescue units are allocated to incidents in the same manner. We present proportions of Monte Carlo simulation results to those of the Greedy Policy by means of box plots. Any value represents the ratio of objective values (total weighted completion times) between the Monte Carlo simulation and the greedy heuristic. The box plots comprise the means (red dash), the quartiles (ends of box), the lowest/highest datum within 1.5 IQR (whiskers), and all outliers (stars). Thus, if both the Monte Carlo simulation and the benchmark deliver the same assignment and schedule, and thus the same objective value, the ratio would be presented as '1.0'. If Monte Carlo performs better than the benchmark, thus the total weighted completion times are lower, the ratios also tend to be lower. 7 different scenarios (10 instances each) have been generated randomly according to the preconditions in Table 2. All Monte Carlo simulations have been aborted after 1,000 iterations to allow for acceptable runtimes in practice. No significant improvements in the results have been identified thereafter. The number of iterations does affect the running time and the resulting quality of the Monte Carlo simulation. Yet, 1,000 iterations of the Monte Carlo simulation were run within 17 minutes for the worst-case scenario (20,200). Results were obtained in 2 minutes for the (10,20) scenario (best-case). Results of the Greedy heuristic have been generated between several seconds (best-case) and 2 minutes (worst-case).

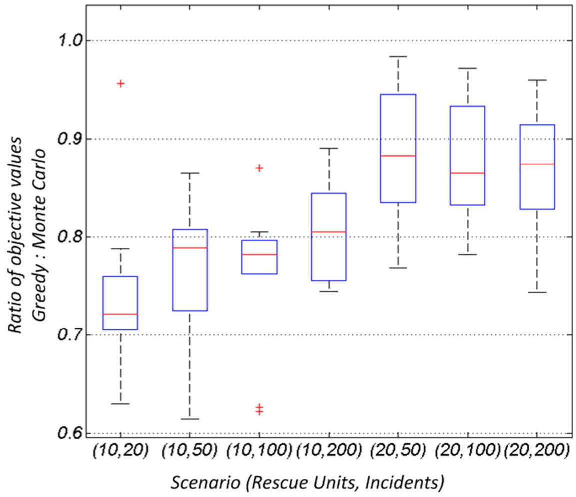

Figure 2. Results indicating the ratio between the heuristics used.

As Figure 2 indicates, the Monte Carlo simulation performs better than the *Greedy Policy*. Ranges of deviation of simulation results were acceptable for all problem scenarios and all results did not exceed the benchmark (proportions ≤ 1.0). In scenario (10,20), the Monte Carlo simulation was even able to generate a total weighted completion time of less than a quarter of which would have been caused by the greedy heuristic. Monte Carlo delivered damage reductions of at least 10%-20% on average compared to the benchmark. Apparently, the ratios are closer to 100% the more complex the scenarios get (starting from 20 rescue units). This phenomenon is not surprising as the fraction of the solution space that gets evaluated by the Monte Carlo simulation declines with increasing instance size. A countermeasure would be to increase the number of iterations in the Monte Carlo simulation, which in turn would require having available more computing power than we had (in order to sustain the efficiency and thus the applicability of the Monte Carlo simulation in practice). Based on the results at hand, we observe a high coefficient of variation for some scenarios that we explain as a consequence of "fuzzifying" the parameters, which may reflect the cost of incorporating linguistic vagueness.

All results were subjected to the Shapiro-Wilk test (Shapiro and Wilk, 1965) to prove normality. Pre-proven normality holds as necessary condition for further analysis: results of significance tests expressed that the simulations of all our models do outperform the benchmark within the confidence intervals of a 95% significance level except for instance (10,100) where a normal distribution of the results was rejected. Our results attest that solving our models with Monte Carlo outperforms the heuristic which is applied in practice.

CONCLUSION

The management of natural disasters poses immense challenges ranging from informational uncertainty to the coordination problem of distributed, heterogeneous rescue units, since disasters continue to hit our societies. Although NDM has evolved to a research discipline where IS artifacts have already been proposed, decision support procedures for assignments and schedules of rescue units have mostly been neglected in research.

Addressing this lack in research, the paper proposes a quantitative decision support model for the optimal allocation of distributed, heterogeneous rescue units based on fuzzy set theory to deal with non-statistical informational uncertainty. Requirements identified in the literature and in interviews are accounted for. Our Monte Carlo-based solution heuristic offers decision support timely to any commander. While the proposed decision model may be particularly useful in the presence of complex situations with large numbers of rescue units and incidents, any assignments and schedules of rescue units determined through computation are not intended to replace the actual decision making process nor the autonomy of commanders but may serve as valuable decision support only. When seen in the context of solving a sequence of model instances over time, our decision support approach also accounts for the dynamics and unpredictability in disaster management situations, which are rooted in the occurrence of new or newly assessed incidents and a changing availability of rescue units. Thereby, our approach can be embedded in a time-continuous optimization process.

Due to the computational hardness of our decision model, we draw on Monte Carlo simulation and computationally demonstrated its benefits. The results show that there is large potential to improve a greedy heuristic to allocate and schedule rescue units. To conclude, we are aware that our research still has some limitations and invites for future work: (1) We exclude the possibility that rescue units may fatigue and thus refrain from a reduction in performance of rescue units over time. (2) Our model does not account for time windows of incidents. Such windows are appropriate when casualties have a finite “time to live” to be rescued. (3) The model does not consider pre-emptive approaches. (4) As real-life data-sets merely exist, all scenarios had to be randomly generated. Thus, empirical research is necessary to gather more realistic data.

REFERENCES

1. Airy, G., Mullen, T., and Yen, J. (2009) Market Based Adaptive Resource Allocation for Distributed Rescue Teams. *Proceedings of the 6th International ISCRAM Conference – Gothenburg, Sweden*.
2. Altay, N., and Green III, W. G. (2006) OR/MS research in disaster operations management. *European Journal of Operational Research*, *175*(1), 475–493.
3. Blazewicz, J., Dror, M., and Weglarz, J. (1991) Mathematical programming formulations for machine scheduling: A survey. *Operations Management Research, 51*(3), 283-300.
4. Buckley, J. J., and Eslami, E. (2002) *An Introduction to Fuzzy Logic and Fuzzy Sets*. Heidelberg: Physica Verlag.
5. Buckley, J. J., and Jowers, L. J. (2008) *Monte Carlo Methods in Fuzzy Optimization*. Berlin, Heidelberg: Springer-Verlag.
6. Chawla, S. (2011) *Japan Earthquake Resources: Aid Organizations, Charities and the Travel Industry Response*. Retrieved from http://www.petergreenberg.com/b/Japan-Earthquake-Resources:-Aid-Organizations,-Charities-and-the-Travel-Industry-Response/-651186966021527775.html, last accessed Dec 08 2011.
7. Chen, R., Sharman, R., Rao, H. R., and Upadhyaya, S. J. (2008) Coordination in emergency response management. *Communications of the ACM*, *51*(5), 66–73.
8. Comes, T., Conrado, C., Hiete, M., Kamermans, M., Pavlin, G., and Wijngaards, N. (2010) An intelligent decision support system for decision making under uncertainty in distributed reasoning frameworks. *Proceedings of the 7th International ISCRAM Conference.* Seattle, USA.
9. Comfort, L. K. (1999) *Shared Risk : Complex Systems In Seismic Response*. Amsterdam: Pergamon.
10. Comfort, L. K., Ko, K., and Zagorecki, A. (2004) Coordination in Rapidly Evolving Disaster Response Systems The Role of Information. *American Behavioral Scientist*, *48*(3), 295–313.
11. Deutsche Presse-Agentur. (2011) *Poor communication holds up aid to Japan after quake and tsunami*. Retrieved from http://www.monstersandcritics.com/news/asiapacific/news/article_1629334.php/Poor-communication-holds-up-aid-to-Japan-after-quake-and-tsunami, last accessed Dec 08 2011.
12. Dmitracova, O. (2010) *Poor coordination biggest problem for relief work - report*. Retrieved from http://www.alertnet.org/db/an_art/60725/2010/01/10-155441-1.htm, last accessed Dec 08 2011.
13. Engelmann, H., and Fiedrich, F. (2007) Decision Support for the Members of an Emergency Operation Centre after an Earthquake. *Proceedings of the 4th International ISCRAM Conference - Delft, The Netherlands*.
14. FEMA (2011) *FEMA Disasters & Maps*. Retrieved from http://www.fema.gov/hazard/, last accessed Dec 08 2011.
15. Fiedrich, F., Gehbauer, F., and Rickers, U. (2000) Optimized resource allocation for emergency response after earthquake disasters. *Safety Science*, *35*(1-3), 41–57.
16. IFRC. *Disaster management - IFRC*. Retrieved from http://www.ifrc.org/en/what-we-do/disaster-management/, last accessed Dec 08 2011.
17. Klingner, B. (2011) *Fukushima Crisis Shows Weakness in Japanese Crisis Management*. Retrieved from http://www.heritage.org/research/commentary/2011/10/fukushima-crisis-shows-weakness-in-japanese-crisis-management, last accessed Dec 08 2011.
18. Klir, G. J., and Yuan, B. (1995) *Fuzzy sets and fuzzy logic: Theory and applications*. Upper Saddle River, N.J.: Prentice Hall PTR.

19. Krolicki, K. (2011) *Special Report: Mistakes, misfortune, meltdown: Japan's quake*. Retrieved from http://www.reuters.com/article/2011/03/17/us-japan-quake-meltdown-specialreport-idUSTRE72G65Z20110317, last accessed Dec 08 2011 .
20. Nair, R., Ito, T., Tambe, M., and Marsella, S. (2002) Task Allocation in the RoboCup Rescue Simulation Domain: A Short Note. *Lecture Notes in Computer Science*, *2002*(2377), 751–754.
21. New South Wales Government. *Emergency Operations*. Retrieved from http://www.ambulance.nsw.gov.au/about-us/Emergency-Operations.html, last accessed Dec 08 2011.
22. Ramchurn, S. D., Rogers, A., Macarthur, K., Farinelli, A., Vytelingum, P., Vetsikas, I., and Jennings, N. R. (2008) Agent-based coordination technologies in disaster management. *Proceedings of the 7th International Joint Conference on Autonomous Agents and Multiagent Systems*, 1651–1652.
23. Reijers, H. A., Jansen-Vullers, M. H., Zur Muehlen, M., and Appl, W. (2007) Workflow management systems + swarm intelligence = dynamic task assignment for emergency management applications. *Proceedings of the 5th International Conference on Business Process Management*, *2007*, 125–140.
24. Rolland, E., Patterson, R., Ward, K., and Dodin, B. (2010) Decision support for disaster management. *Operations Management Research*, *3*(1), 68–79.
25. Sanders, S. (2011) Japan’s Sendai earthquake: One month later. *The Washington Post*. Retrieved from http://www.washingtonpost.com/blogs/blogpost/post/japans-sendai-earthquake-one-month-later/2011/04/11/AFdrZtLD_blog.html, last accessed Dec 08 2011.
26. Shapiro, S. S., and Wilk, M. B. (1965) An Analysis of Variance Test for Normality (Complete Samples). *Biometrika*, *52*(3/4), 591–611.
27. U.S. Department of Homeland Security (2008) *Homeland Security Advisory System--Guidance for Federal Departments and Agencies*. Retrieved from https://www.dhs.gov/files/programs/gc_1156876241477.shtm, last accessed Dec 08 2011.
28. Wex, F., Schryen, G., and Neumann, D. (2011) Intelligent Decision Support for Centralized Coordination during Emergency Response. *Proceedings of the 8th International ISCRAM Conference.* Lisbon, Portugal.
29. Zadeh, L. A. (1965) Fuzzy sets. *Information and Control*, *8*(3), 338–353.
30. Zimmermann, H. J. (2000) An application-oriented view of modeling uncertainty. *European Journal of Operational Research*, *122*(2), 190–198.

APPENDIX

Proof of NP-hardness

The Fuzzy Decision Model (M1) is a generalization of the machine scheduling problem “Identical parallel machine non-preemptive scheduling with minimization of sum of completion times” (M2), which is NP-hard (Blazewicz et al., 1991): if we map incidents on jobs and rescue units on machines, then the generalization refers to the fact that our problem provides for setup times (travel times), non-identical machines, and constraints on the assignment of rescue units to incidents. Given an instance of M2, we can map this instance onto an instance of M1 (in polynomial time) by ignoring each parameter that belongs to a fuzzy set, by setting $\widetilde{s_{ij}^k} = 0$ for all jobs i,j and for all machines k, by setting $\widetilde{p_i^{k_1}} = \widetilde{p_i^{k_2}}$ for all jobs i and all machines k_1 and k_2, and by setting $cap_{kl} = 1$ for all rescue units k and for all incidents i. Thus, our problem is NP-hard, too. Integrating Fuzzy Set Theory in this proof even raises the complexity.