On the Exploitation of Automated Planning for Efficient Decision Making in Road Traffic Accident Management

Lukáš Chrpa and Mauro Vallati

Abstract— **Automated Planning can be fruitfully exploited as a Decision Support toolkit that, given a specification of available actions (elementary decisions to be taken), an initial situation and goals to be achieved, generates a plan that represents a (partially ordered) sequence of such elementary decisions that once performed the required goals are achieved. Road Traffic Accident Management is a life-critical task that deals with effective planning of emergency response when accidents occur, in order to mitigate negative effects, especially saving human lives that might be in imminent danger.**

In this paper, we exploit Automated Planning in the Road Traffic Accident Management domain. We specifically focus on providing necessary treatment for victims injured during accidents. This involves coordination of medical teams responsible for providing medical treatment to the victims and fire brigades that are required to release victims trapped in damaged vehicles. An empirical analysis, based in the region of West Yorkshire (UK) with a number of real accidents recently occurred there, shows the suitability of the proposed Automated Planning approach to be used in time-critical conditions, and confirms the effectiveness of the generated plans. We also demonstrated its usefulness as a tool for evaluating the impact of additional resources, in order to provide guidance for future investments.

I. INTRODUCTION

Managing emergency response to traffic accidents is crucial to mitigate their consequences, especially to save human lives that might be in danger after an accident occurs, and to reduce the economical impact on the society. Effective managing of emergency response involves coordinating medical staff that is responsible for providing first-aid to accident victims, ambulances that are necessary for transporting seriously injured victims to hospitals, and fire brigades that have to assist in cases where a victim is trapped in damaged vehicle. Traffic accident management is also subject to specific local regulations [1] with strictly defined (medical) response times [2] and procedures. This provides a challenge for accident management coordinators because they have to take a number of critical decisions in a very short time. Currently, most of these decisions are taken by humans. However, humans in charge of taking decisions are usually under a huge pressure –particularly in case of multiple accidents– and must react quickly: chances of making errors are therefore high.

Remarkably, there has been work in the area of autonomous systems for supporting human experts. However, it is primarily focused to the Search and Rescue domain [3], [4], mainly thanks to the RoboCup Rescue Robot and Simulation competitions [5]. In emergency response, the research focuses on the problem of determining a (nearly) optimal coverage of emergency services [6] where various techniques such as genetic programming [7] or fuzzy reasoning [8] have been used. Models predicting the likeliness of medical incident occurrence, which can assist the emergency response controllers in their decision making, have been developed [9]. For simulating ambulance deployment in urban areas several systems have been developed [10], [11]. However, the aforementioned approaches give a little assistance in decision making for emergency response controllers.

Automated planning, which deals with the problem of finding a plan (a sequence of actions) that transforms the environment from an initial state to some desired goal state [12], is an effective tool for decision making. Plans consist of information about which action and at which time it has to be executed, in order to achieve given goals. In traffic accident management, for example, a plan embodies a procedure which if correctly followed will ensure that the all accident victims have received appropriate treatment while taking into account constraints (e.g. the number of medical staff). Moreover, with numerous generic planning engines that accept the description of planning problems in a standard language such as PDDL [13], it is easy to apply Automated Planning as part of larger intelligent systems (e.g., the recent work in marine robotics [14]).

In this paper, we aim to use Automated Planning as an effective Decision Support tool in the Road Traffic Accident Management domain. Automated Planning provides a powerful toolkit for decision support that has already been used in real-world applications, including Urban Traffic Control [15], [16]. Applying Automated Planning for decision support in Traffic Accident Management has been considered by Özbay et al. [17], where probabilistic models have been used for planning operations –but not for reacting to actual reported accidents–, and by Shah et al. [18], which considered "classical" domain-independent planning. It should be noted that in the work of Shah et al., the traffic accident management domain was mainly investigated as a case study for comparing Knowledge Engineering techniques; one of the resulting domain models was included in the temporal track of the last International Planning Competition [19]. In terms of modelling of real-world scenarios, it was quite simplistic. Recently, the work has been also adapted for general incident management [20].

Inspired by the work of Shah et al. [18], here we specify and develop a planning domain model that complies with standards provided by the National Health Services (NHS) in the UK [2]. The domain model is encoded in

L. Chrpa and M. Vallati are with School of Computing & Engineering, University of Huddersfield, HD1 3DH Queensgate, Huddersfield, UK

Planning Domain Definition Language (PDDL) [13] that is widely supported by the large number of existing domain-independent planning engines. Our approach is empirically evaluated on several scenarios that involve a part of the area of West Yorkshire (UK) with actual locations of emergency services and frequent traffic accident spots. The results are thoroughly discussed in order to understand strengths and drawbacks of using Automated Planning in Traffic Accident Management as well as to indicate promising avenues for future research.

II. Background

Automated Planning can deal with different levels of expressiveness. In this work we focus on *Temporal Planning*, which is a subclass of Automated Planning that reasons with actions whose execution takes time (so called "durative actions") in a deterministic and fully observable environment. The environment is described by first-order logic *predicates* and *numeric fluents*. *Actions* are specified via their execution *duration*, *preconditions* which are logical expressions that must hold in order to make the action executable, and *effects* which are sets of literals or fluent assignments that take place when the action is executed. In PDDL 2.1, preconditions can take place just before the action is executed, during execution of the action and just before finishing execution of the action. Similarly, effects can take place just after stating execution of the action, or just after finishing execution of the action [13]. A *Planning Domain Model* consists of predicates, numeric fluents and actions. A *Planning Problem Description* consists of a set of objects, an initial state (a set of grounded predicates and fluent assignments), and a set of goals (logical expressions). A *plan* is a set of pairs in the form of $\langle$timestamp,action$\rangle$ such that executing these actions in the corresponding timestamps (it must always be possible) transforms the environment from the initial state to some state where all the goals are satisfied (i.e. a goal state).

III. Problem Specification

Traffic accidents are events that put human lives in danger and cause huge economic costs, so it is crucial to respond quickly and efficiently to mitigate their consequences. To do so, it is necessary to coordinate teams of professionals (e.g. paramedics or fire brigades) as well as manage limited amount of resources (e.g. ambulances). Our primary focus is directed towards rescuing and providing necessary medical treatment to accident victims to maximize chances of their survival.

When an accident occurs and is reported, information about victims on the site of the accident that need to be treated are given. According to NHS guidelines [2], patients that need assistance are prioritised following a 3-categories (A, B and C) schema. Category A includes patients in life-threatening conditions, category B are patients in serious but not life-threatening conditions, while category C includes not serious patients. Patients in A and B categories should be reached within very strict time limits, respectively 8 and 19 minutes from the accident report. NHS also uses the reaching time as metrics for evaluating the local services. Nationally, the target is to reach the 75% of category A patients within the 8 minutes limit, and the 95% of category B patients within the 19 minutes limit.

In this work, we consider patients in categories A and B, as well as those who are trapped in damaged vehicles. Notice that patients in category C are not of interest in our case, given the not strict (up to 1 hour) time limits for treating them. The goal is to have all the victims given first-aid to their injuries and been untrapped (if they were initially trapped in damaged vehicles), and the seriously injured victims been transported into hospitals by ambulances. The plan achieving such a goal provides the emergency response coordinators a procedure that distributes tasks to teams of medics and fire brigades, and provides an allocation of available resources which are required for the execution of the tasks (e.g. ambulances).

A. Domain Model Specification

Given the "high-level" problem specification, one might have an impression about actions that have to be taken in order to achieve the goal (rescue and treat traffic accident victims). To be precise, the conceptualisation of the problem requires to specify *object types*, which stand for classes of objects considered in the planning process, *predicates and numeric fluents*, which describe the environment, and *actions*, which modify the environment [18].

Object types in our model are divided into four main categories: *Assets*, *Agents*, *Locations* and *Victims*. Assets are further divided into *Static Assets*, which consist of hospitals and fire stations, and *Mobile Assets*, which consist of fire brigades, ambulances, medical cars and medical motorbikes. Agents are human professionals, in particular drivers, firemen, medics (paramedics and emergency medical technicians) that carry out "role-specific" tasks (e.g. giving first-aid).

The static part of the environment, i.e., aspects that do not change during execution of the plan, is related to the road network and some properties of the assets. We consider an abstract road network, where only locations of interest (e.g. asset or accident locations) are considered and existence of a road between two locations is denoted by a predicate *connected* and the distance between connected locations is denoted by a numeric fluent *distance*. Average speed of a mobile asset is captured by a numeric fluent *speed*. Note that although we assume that average speed of an asset does not change during the planning episode, it might be initially set differently for different planning episodes (e.g. average speed is smaller for rush hours). We provide restrictions on whether an agent can enter/board an asset, which is captured by a predicate *compatible*. For mobile assets that need a "driver agent" in order to move, we denote such a fact by a predicate *need-driver*. For static assets, their locations are represented by a predicate *at*.

The dynamic part of the environment considers aspects that can be modified by executing planned actions. Mobile assets, agents and victims can change their locations

```
(:durative-action load-victim
:parameters (?v - victim ?x - ambulance
?m - medic ?l - location)
:duration (= ?duration 5)
:condition (and
   (at start (at ?v ?l))
   (at start (free ?x))
   (over all (aided ?v))
   (over all (at ?x ?l))
   (over all (at ?m ?l))
   (at start (available ?m)) )
:effect (and
   (at start (not (at ?v ?l)))
   (at start (not (free ?x)))
   (at start (not (available ?m)))
   (at end (in ?v ?x))
   (at end (available ?m)))
)
```

Fig. 1. PDDL encoding of the `load-victim` operator.

– their current location in a given timestamp is denoted by a predicate *at* (similarly to the static assets). Agents, in addition, can be inside assets (denoted by a predicate *inside*). Victims, similarly to agents, can be in ambulances or hospitals (denoted by a predicate *in*). Clearly, state invariants apply. A mobile asset can be at at most one location at time. An agent (victim) can be either at one location or inside (in) one asset at time. An ambulance (a type of mobile asset) can be *free* or carry at most one victim (i.e., the victim is *in* the ambulance). Notice that capacity of mobile assets in terms of how many agents it can accommodate can be ensured by the *compatible* predicates. That is that if the number of instances of the *compatible* predicate with respect to a given asset is not higher than its capacity, then there is no plan in which the capacity can be exceeded. Victims can be either *slightly injured* (corresponding to the NHS category B), *seriously injured* (category A), or *first-aided*. Also, victims can be either *trapped* or *untrapped*. However, for clear reasons, it is not possible that a victim is both *trapped* and *first-aided*.

In temporal planning, action execution takes some time. An agent can perform (or be involved in) at most one action at time, such constraint has been encoded using well-known PDDL techniques, and is therefore omitted from the following description of actions. Let t_s be the timestamp at which an action is executed, and t_e be the timestamp at which an action execution finishes. We have defined the following actions:

- **move(x,a,l$_1$,l$_2$)** – moves a mobile asset x driven by an agent a from a location l_1 to a location l_2. As a precondition $at(x, l_1)$ must hold in t_s, $connected(l_1, l_2)$, $inside(a, x)$ as well as that a is a driver if *need-driver*(x) is true must hold in $[t_s, t_e]$ (over the whole execution interval). As an effect $at(x, l_1)$ becomes false in t_s and $at(x, l_2)$ becomes true in t_e. The action duration is determined from $distance(l_1, l_2)$ and $speed(x)$.
- **board(x,a,l)** – an agent a boards/enters an asset x at a location l. As a precondition $at(a, l)$, $at(x, l)$ as well as $compatible(a, x)$ if x is a mobile asset must hold in t_s. As an effect $at(a, l)$ becomes false in t_s and $inside(a, x)$ becomes true in t_e.
- **debark(x,a,l)** – an agent a debarks/exits an asset x at a location l. As a precondition $inside(a, x)$, $at(x, l)$ must hold in t_s. As an effect $inside(a, x)$ becomes false in t_s and $at(a, l)$ becomes true in t_e.
- **untrap(f,v,l)** – a fireman f untraps a victim v at a location l. As a precondition $at(f, l)$, $at(v, l)$ must hold in $[t_s, t_e]$ and $trapped(v)$ must hold in t_s. As an effect $trapped(v)$ becomes false at t_e.
- **firstaid(m,v,l)** – a medic m gives a first-aid to a victim v at a location l. As a precondition $at(m, l)$, $at(v, l)$ must hold and $trapped(v)$ must not hold in $[t_s, t_e]$ and *slightly-injured*(v) or *seriously-injured*(v) must hold in t_s. If *seriously-injured*(v) holds in t_s, m must be a paramedic. As an effect *slightly-injured*(v) or *seriously-injured*(v) becomes false, and $aided(v)$ becomes true at t_e. The action duration is larger for the *seriously-injured*(v) case.
- **load-victim(v,x,m,l)** – A medic m loads a victim v into an ambulance x at a location l. As a precondition $at(m, l)$, $at(x, l)$ and $aided(v)$ must hold in $[t_s, t_e]$ and $at(v, l)$ and $free(x)$ must hold in t_s. As an effect $at(v, l)$ and $free(x)$ become false in t_s, and $in(v, x)$ becomes true in t_e.
- **hospitalise(v,m,x,h,l)** – A medic m hospitalises a victim v from an ambulance x to a hospital h at a location l. As a precondition $at(x, l)$ and $at(h, l)$ must hold in $[t_s, t_e]$, and $in(v, x)$, $at(m, l)$ must hold in t_s. As an effect $at(m, l)$ and $in(v, x)$ become false in t_s, and $in(v, h)$, $inside(m, h)$ and $free(x)$ become true in t_e.

Figure 1 shows the PDDL encoding of the described **load-victim(v,x,m,l)** operator. It is worth noting the presence of the $available$ predicate, which is used for ensuring that an agent is performing at most one action at time.

IV. EVALUATION OF THE APPROACH

This section is devoted to evaluate the ability of the proposed approach in managing limited resources for effectively handling road traffic accidents.

A. Settings

The Planning domain model and problems have been encoded in PDDL 2.1. For solving such problems, we selected Yahsp3 [21], a planner that achieved remarkable results in recent International Planning Competitions (IPCs) [19] –the major competition run by the planning community– and that is able to handle the set of PDDL features required by the encoded domain model. It should be noted that Yahsp3 is not an optimal solver –i.e. plans identified are not optimal with regard to a considered metric– but it tends to be very quick, and to provide good quality plans. Furthermore, we also considered the LPG planner [22], due to its good performance in previous editions of the IPC. Both systems can be used as *anytime* solvers, i.e. they are able to improve the quality of generated plans if additional CPU-time is given. In this context, quality is measured in terms of overall time required to achieve the specified goals.

The planners have been run on a system equipped with 2.5 Ghz Intel Core 2 Quad Processors, 4 GB of RAM and Linux operating system. The generated plans have been validated using the well-known VAL tool [23]. This was done in order to check their correctness with regards to the designed domain model, and also to identify the presence of bugs or flaws in the model.

Our analysis has been focused on the region controlled by the Yorkshire Ambulance Service (YAS) and by the West Yorkshire Fire Service. YAS is in charge on the whole Yorkshire county, which is located in the north-east part of United Kingdom. However, such large area is divided into several regions. Here, following the way in which the YAS service is organised, we considered the Bradford, Calderdale and Kirklees regions (BCK). In the mentioned region there are 7 ambulance stations –where ambulances and medical equipment are stored, ready to be used– and 5 hospitals, assets that are able to receive emergency patients. In order to simulate a busy condition, fleet and crews have been distributed as follows: each hospital has one ambulance and the corresponding crew available; each hospital has either a car or a motorbike available, with an associated paramedic. Beside the ambulance service resources –which still represent the main focus of our investigation– two fire stations, which are the main stations in the BCK area, have also been considered. Each fire station has a pump/ladder vehicle and the corresponding crew available. Figure 2 (coloured) shows an overview of the considered region.

In terms of scenarios used in this evaluation, we simulated two accidents recently happened, involving respectively 2 and 6 victims[1]. Moreover, we also included an accident at Ainley top (located between Huddersfield and Brighouse, see Figure 2), which simulates the fact that the huge traffic flow on M62 –the main highway in the area– leads to frequent incidents. In total, 6 category A (seriously injured) and 4 category B (slightly injured) patients (victims) are involved in the three car accidents.

B. Effective Emergency Response for Victims in Life-threatening Conditions

As previously mentioned, NHS poses strict time targets for reaching life-threatening and seriously injured victims. It is therefore critical that the proposed planning-based approach is able to efficiently and effectively manage the available resources, in terms of crews and vehicles, to meet the nationally set targets.

In a first set of experiments, designed for testing the fact that the proposed system is effectively able to manage and exploit available resources, we considered a scenario in which the presence of a single victim in life-threatening condition is reported, and it should be first-aided as soon as possible. We run 10 different experiments by changing the position of the victim: s/he was placed respectively at the three previously mentioned accident places, and in other random locations of the considered region. It is worthy reminding that a victim in life-threatening conditions must be reached by a paramedic for the required care.

[1]The interested reader can find more information about the mentioned accidents here: http://goo.gl/vctsHZ, http://goo.gl/Ag5ePP

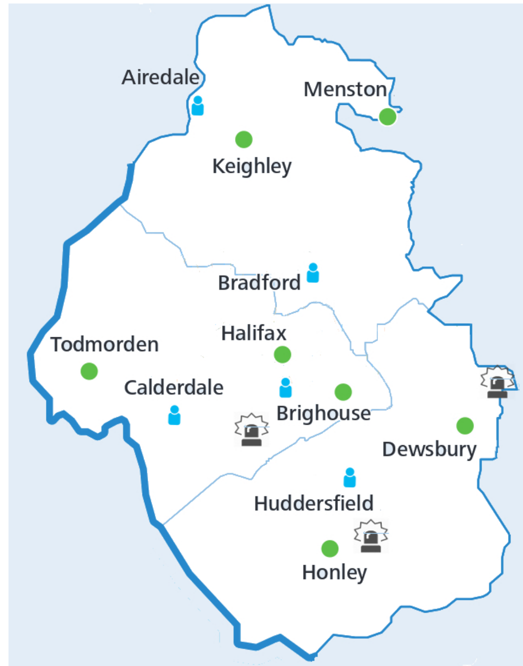

Fig. 2. The region considered in the experimental analysis: this includes the Calderdale, Bradford and Kirklees councils, within the Yorkshire county of United Kingdom. Green dots indicate the position of ambulance stations; hospital locations are represented by the blue patient icon. Emergency icons are used for indicating the places of accidents as considered in our scenarios.

Remarkably, in every plan generated by the considered planning engines for the described scenarios, the available paramedic which is able to reach the victim in the shortest time –either because closer to the accident location, or using a faster vehicle– is always selected to succour the victim. This allows to fulfil the national NHS target for category A patients in all the cases in which it is actually possible to reach the accident locations in less than 8 minutes.

Clearly, victims in life-threatening conditions must also be hospitalised as soon as possible: the support of a paramedic at the accident location is fundamental, but not sufficient, to maximise their chance of survival. For this reason we run a second set of experiments, considering the same previously described scenarios, but requiring that victims are also hospitalised. Also in such cases we observed that the plans generated by Yahsp3 and LPG are very valuable. Whenever possible, a paramedic is sent to the accident location for providing first aid and, in the meanwhile, an ambulance is sent for transporting the victim to the closest hospital. When both ambulances and other vehicles take almost the same time to reach the accident place, in some cases planners decide to send only an ambulance: the delay in providing first-aid is negligible (usually less than a minute).

Finally, in order to evaluate the coordination between

```
5: (debark Paramedic RRV-Brighouse
AinleyTop)
10.3: (debark FireCrew FireVehicleHalifax
AinleyTop)
10.3: (untrap FireCrew Victim1 AinleyTop)
20.3: (firstaid Paramedic Victim1
AinleyTop)
```

Fig. 3. Part of plan (and timestamp at which actions start) provided by Yahsp3 involving both fire brigades and paramedics. RRV stands for Rapid Response Vehicles, i.e. a car used by paramedics.

ambulance services and fire brigades achievable by using our planning approach, we considered cases in which the category A victim is trapped into a vehicle, at the different accident locations. In such scenario, planners correctly decide to send the closest available fire-fighting vehicle –with the appropriate crew–. In the meanwhile, a paramedic is sent to the accident location, for providing support. Figure 3 shows part of an example plan. In the shown case, the rapid response vehicle arrived before the fire brigade at the accident location, and the paramedic can start supporting the victim as soon as s/he has been untrapped.

In terms of required runtime, the planners are able to generate plans for handling all the scenario considered in this set of experiments very quickly; less than 0.5 CPU-time seconds are necessary.

C. Response to Multiple Concurrent Accidents

In the previous section, we focused our analysis on cases in which a single victim required support. However, significantly more complex –and to some extent more realistic– scenarios are those in which more than one accident is reported, and need to be handled, at the same time. Here we therefore assess the performance of the proposed planning-based approach in handling the three accidents (located accordingly to Figure 2) concurrently, the goal is that all the category A victims are hospitalised, and category B are first-aided, as soon as possible. We emphasise that the number of available vehicles is extremely limited, with regards to the number of critical victims in the need of support: five ambulances are available, while ten victims are reported. Moreover, a few medical vehicles are located at premises that are too far from the accident locations to be useful.

As a first remark, we note that despite the complexity of the modelled problem, both LPG and Yahsp3 are able to generate a valid plan using very short runtimes: less than 1 CPU-time second is required by YAHSP, while LPG usually takes between 1 and 3 seconds (it is a randomised planner, so CPU-time can differ between runs). However, the quality of first generated plans –in terms of temporal length of the plan– can be low. For this reason we let planners run for 10 CPU-time seconds, in order to allow them improving such plans. Even in the time-critical domain of emergency management, 10 seconds is a reasonable amount of time that can be invested for obtaining good quality plans. Table I shows a comparison of the performance of the two planning

	LPG	Yahsp3
Runtime (seconds)	1.2 – 3.8	0.7
Plan Quality (minutes)	143 – 183	131
1st Aided Cat. A (minutes)	8 – 28	9
1st Hospitalised Cat.A (minutes)	52 – 75	26
1st Aided Cat. B (minutes)	8 – 38	8

TABLE I

RUNTIME (CPU-TIME SECONDS) NEEDED BY LPG AND YAHSP3 FOR PROVIDING THE FIRST SATISFICING PLAN FOR HANDLING THREE CONCURRENT CAR ACCIDENTS. QUALITY OF GENERATED PLANS IS EVALUATED IN TERMS OF EXECUTION DURATION. INFORMATION ABOUT THE TIME NEEDED TO: AID THE FIRST CATEGORY A VICTIM, HOSPITALISE THE FIRST CATEGORY A VICTIM, AND AID THE FIRST CATEGORY B VICTIM IS PROVIDED. LPG IS A RANDOMISED PLANNER: BEST AND WORST RESULTS OBTAINED ARE SHOWN.

engines. Performance is shown in terms of overall temporal length of plans, time needed to reach the first category A and category B victims, and the time needed for transporting to hospital the first victim among those involved in the three accidents. By analysing the plans, we observed that all the vehicles and crews close to accident locations are effectively exploited by the planners.

D. Assess the Usefulness of Additional Resources

Noteworthy, the planning-based approach proposed in this paper can also be exploited as a tool for assessing the usefulness of additional resources with regards to the already available vehicles fleet and personnel, in order to provide some guidance on possible future investments. Intuitively, the evaluation of usefulness of additional resources can be done by testing their impact on plans generated for handling expected recurrent situations.

As a case study, we aim at evaluating if it would be worthy to extend the considered YAS fleet with an additional ambulance, and the corresponding crew, and where it would be better to place it for maximising usefulness. We consider our three accidents scenario as a typical recurrent situation that has to be faced by the Yorkshire Ambulance Service. The investigation is done by positioning the additional ambulance and crew in each of the YAS premises, generating and solving the corresponding planning problem, and measuring the impact on the plan's quality. By using a 10 seconds CPU-time limit, the overall evaluation takes at most 130 seconds.

According to the performed analysis, the additional ambulance and crew should be placed at Dewsbury in order to maximise their impact for handling the modelled scenario. Specifically, their exploitation would reduce the length of the best plan by approximately 10%: from 131 to 117 minutes. This also indicates that most of the victims in life-threatening conditions are hospitalised earlier, with a significant positive repercussion on their survival chances. On the other hand, it has been observed that positioning the additional resources in the Airedale or Menston YAS premises will have no effect on plans for the studied scenario. In other words, they will not be exploited.

V. DISCUSSION AND CONCLUSION

Effective management of emergency response to traffic accidents is crucial for maximising the chances of survival of victims. Remarkably, such management requires a very high level of coordination between a number of entities, like medical staff, medical vehicles and fire brigades. With the aim of providing useful guidelines and improving the quality of services, clear and strict regulations and guidelines have been defined and put in place by national bodies, such as the National Health Services (NHS) in the UK. However, despite its importance and complexity, the management of emergency response task is currently mainly performed by human experts.

In order to assist experts in the stressful and time-critical duty of road traffic accident management, here we investigated the exploitation of automated planning techniques as an effective decision support tool. The main contribution of our work are: (i) a PDDL 2.1 encoding of the traffic accidents management domain, which embeds guidelines from the NHS regulations; (ii) a thorough experimental analysis, that considers the actual premises of NHS in West Yorkshire, and allows to highlight strengths and weaknesses of the proposed approach; and (iii) a comparison between plans and strategies generated by two state-of-the-art domain-independent planning engines.

The performed experimental analysis demonstrates the extent to which the proposed approach is able to efficiently and effectively coordinate available resources for emergency response to traffic accidents. Remarkably, plans are generated very quickly –in a few seconds–, and this demonstrates the potential of the approach for being used in time-critical conditions. Moreover, an analysis of generated plans confirmed their overall quality feasibility: they can therefore provide useful suggestions for the human experts in charge of managing the emergency.

Finally, it is worth noting that the proposed PDDL 2.1 encoding can be easily extended. For instance, an extended model can also consider different sort of vehicles, such as helicopters. Moreover, the knowledge about the scenario can be easily improved and updated: traffic conditions can be encoded, and problems can be modified if new pieces of information are made available to the emergency operator.

For the future, we plan to integrate our approach into some of the existing simulation platforms. This will allow us to see planning and plan execution from accident coordinator's point of view which will validate the concept for a potential use in practice. We are also interested in considering different, and more expressive, languages for describing the considered problem. Finally, we are interested in evaluating the exploitation of decentralised (multi-agent) strategies for stretching the importance of coordination between the different agents involved in handling traffic accidents emergencies.

REFERENCES

[1] Benesch, *Traffic Incident Management Operations Guidelines*, Iowa Department of Transportation, Federal Highway Administration, Washington, USA, 2011.

[2] *The Health and Social Care Act 2008 (Regulated Activities)*, 781st ed., National Health Service, 2010.

[3] P. Doherty and F. Heintz, "A delegation-based cooperative robotic framework," in *2011 IEEE International Conference on Robotics and Biomimetics, ROBIO 2011, Karon Beach, Thailand, December 7-11, 2011*, 2011, pp. 2955–2962.

[4] J. George, P. B. Sujit, and J. B. Sousa, "Search strategies for multiple UAV search and destroy missions," *Journal of Intelligent and Robotic Systems*, vol. 61, no. 1-4, pp. 355–367, 2011.

[5] H. L. Akin, N. Ito, A. Jacoff, A. Kleiner, J. Pellenz, and A. Visser, "Robocup rescue robot and simulation leagues," *AI Magazine*, vol. 34, no. 1, pp. 78–86, 2013.

[6] M. Zhao and Q. Chen, "Risk-based optimization of emergency rescue facilities locations for large-scale environmental accidents to improve urban public safety," *Natural Hazards*, vol. 75, no. 1, pp. 163–189, 2015.

[7] N. Geroliminis, K. Kepaptsoglou, and M. G. Karlaftis, "A hybrid hypercube - genetic algorithm approach for deploying many emergency response mobile units in an urban network," *European Journal of Operational Research*, vol. 210, no. 2, pp. 287–300, 2011.

[8] C. Araz, H. Selim, and I. Ozkarahan, "A fuzzy multi-objective covering-based vehicle location model for emergency services," *Computers & OR*, vol. 34, no. 3, pp. 705–726, 2007.

[9] H. Setzler, C. Saydam, and S. Park, "EMS call volume predictions: A comparative study," *Computers & OR*, vol. 36, no. 6, pp. 1843–1851, 2009.

[10] S. G. Henderson and A. J. Mason, "Ambulance service planning: simulation and data visualisation," in *Operations research and health care*. Springer, 2005, pp. 77–102.

[11] P. M. S. Silva and L. R. Pinto, "Emergency medical systems analysis by simulation and optimization," in *Proceedings of the 2010 Winter Simulation Conference, WSC*, 2010, pp. 2422–2432.

[12] M. Ghallab and D. Nau and P. Traverso, *Automated Planning Theory and Practice*. Elsevier Science, 2004.

[13] M. Fox and D. Long, "Pddl2. 1: An extension to pddl for expressing temporal planning domains." *J. Artif. Intell. Res.(JAIR)*, vol. 20, pp. 61–124, 2003.

[14] L. Chrpa, J. Pinto, M. A. Ribeiro, F. Py, J. B. de Sousa, and K. Rajan, "On mixed-initiative planning and control for autonomous underwater vehicles," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*, 2015, pp. 1685–1690.

[15] I. Cenamor, L. Chrpa, F. Jimoh, T. L. McCluskey, and M. Vallati, "Planning & scheduling applications in urban traffic management," in *Proceedings of PlanSIG 2014, Annual Workshop of UK Planning & Scheduling Special Interest Group*, 2014.

[16] M. Vallati, D. Magazzeni, B. De Schutter, L. Chrpa, and T. L. McCluskey, "Efficient macroscopic urban traffic models for reducing congestion: a pddl+ planning approach," in *The Thirtieth AAAI Conference on Artificial Intelligence (AAAI)*, 2016.

[17] K. Ozbay, C. Iyigun, M. Baykal-Gursoy, and W. Xiao, "Probabilistic programming models for traffic incident management operations planning," *Annals OR*, vol. 203, no. 1, pp. 389–406, 2013.

[18] M. M. Shah, L. Chrpa, D. Kitchin, T. L. McCluskey, and M. Vallati, "Exploring knowledge engineering strategies in designing and modelling a road traffic accident management domain," in *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence (IJCAI)*. AAAI Press, 2013, pp. 2373–2379.

[19] M. Vallati, L. Chrpa, M. Grzes, T. L. McCluskey, M. Roberts, and S. Sanner, "The 2014 international planning competition: Progress and trends," *AI Magazine*, vol. 36, no. 3, pp. 90–98, 2015.

[20] L. Chrpa and K. Thorisson, "On applicability of automated planning for incident management," in *The international Scheduling and Planning Applications woRKshop (SPARK 2015)*, 2015.

[21] V. Vidal, "Yahsp3 and yahsp3-mt in the 8th international planning competition," in *Proceedings of the 8th International Planning Competition (IPC-2014)*, 2014.

[22] A. Gerevini, A. Saetti, and I. Serina, "Planning through stochastic local search and temporal action graphs in lpg." *Journal of Artificial Intelligence Research (JAIR)*, vol. 20, pp. 239–290, 2003.

[23] R. Howey, D. Long, and M. Fox, "VAL: automatic plan validation, continuous effects and mixed initiative planning using PDDL," in *16th IEEE International Conference on Tools with Artificial Intelligence (ICTAI)*, 2004, pp. 294–301.