

AI Planning Applied to GIS-based Disaster Response

Charles Petrie
Fraunhofer FOKUS Berlin
(visiting scholar)
Berlin, Germany
petrie@cdr.stanford.edu

Agnès Voisard
Agnes.Voisard@fu-berlin.de
Institute of Computer Science
Freie Universität Berlin
Berlin-Dahlem, Germany

ABSTRACT

GIS information is clearly critical for disaster response planning. Formal planning can eliminate otherwise unanticipated conflicts in disaster response planning, particularly in distributed planning. Such conflicts may result in execution-time discovery and resolution, with sub-optimal results. This is especially true in conflicts based upon GIS information as the cases may be subtle, typically resulting from lack of resource access. However, GIS systems do not produce the kind of quantitative information needed for an interface with many other systems, especially planning. New technologies such as GeoSPARQL provide some hope but much work remains. This is a fertile area for research and development.

CCS CONCEPTS

• Computing methodologies → Planning and scheduling.

KEYWORDS

planning, disaster recovery, GIS, constraint satisfaction

ACM Reference Format:

Charles Petrie and Agnès Voisard. 2019. AI Planning Applied to GIS-based Disaster Response. In *5th ACM SIGSPATIAL International Workshop Emergency Management Using GIS (EMGIS'19)*, November 5, 2019, Chicago, IL, USA. ACM, New York, NY, USA, 4 pages. <https://doi.org/10.1145/3356998.3365761>

1 INTRODUCTION

Disasters that befall nations today are becoming both more complex and frequent. Sooner or later the scope of some disaster will outstrip the abilities of people to effectively handle the crisis without computer augmentation. Many disasters require the coordination of distributed plans that are responses to the disaster. The coordination of many people working on different parts of the problem in real time may be poor, with some decisions conflicting with others. These conflicts may not be perceived by humans under stress and with short time-frames until too late to recover gracefully. Property damage and even deaths may occur that were unnecessary. Computer technology to handle such symbolic planning problems has long been available, but does not handle GIS (Geographical Information System) data, which adds complications. Much of the

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

EMGIS'19, November 5, 2019, Chicago, IL, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6965-7/19/11...\$15.00

<https://doi.org/10.1145/3356998.3365761>

technology available for disaster responses concentrates on alerts and special cases. We show that more can, and needs, to be done.

The authors have been working on a project called SANCTUM¹ to show that a particular model of disaster management could be computerized, and what kind of help it might provide, including useful graphics. This paper grew out of our particular approach to showing feasibility.

This paper is organized as follows. In Section 2, we recap the SANCTUM work for context in disaster response planning. In Section 3, we give an example motivating problem that grew out of this work. In Section 4, we discuss formal planning technology and why it is applicable. In Section 4, we discuss the additional complications of information from a GIS and how we might handle it in the next phase of the project, with technology now available, pointing out unresolved important problems. In Section 5, we discuss the additional complications of information from a GIS and how we might handle it in the next phase of the project, with technology now available, pointing out unresolved important problems.

2 SANCTUM PROBLEM SOLVING CONTEXT

In the SANCTUM project, we work to show how to formalize an extensive model of disaster management, focusing upon disaster responses. This model assumes a distributed mode of problem solving now used by the French government. When a disaster is anticipated, a *central crisis management group*, let us call this CCMG, calculates the likely scenarios, warns the affected agencies, and asks for contingency plans.

Those notified would typically be the Prefects of a Department with population centers likely to be affected, the regional transportation managers (trains, road-based vehicles, water-based transport, and airplanes), the regional utility managers, and the managers of large industries. Notified of the kind of danger and time-frame, they would return with general plans to be worked out by the CCMG.

The case study we looked at was the imminent rupture of a large dam: the Barrage de Grand' Maison. We advised that for every such kind of disaster, we could build a meta-plan. In the case of a dam rupture, the procedure would be 1) obtain the names of the valleys located below the dam using an external GIS, 2) obtain the agencies in the valleys that need to be notified from internal databases, 3) use external programs to project the time-to-flood for the assets belonging to the agencies, and then 3) alert them and request their existing contingency plans.

For this alert (versus response) planning, we have prototyped a web service planner using [8] to do this supposing that we could

¹<https://www.cerema.fr/fr/actualites/sanctum-projet-recherche-destine-ameliorer-efficacite>

use some form of web services to access a GIS, a list of agencies and assets, and a time-to-flood prediction program²

Already we are assuming that some knowledge interchange issues are solvable. For instance, we assume that, in some form, we can query the GIS and obtain information such as the assertions ("Valley" "Grand Maison Barrage" "La Romanche") and ("Valley" "Grand Maison Barrage" "Lac du Verney") . This allows the alert planner to branch in creating a particular plan, from the meta-plan, querying the agencies and assets database for each valley. Similarly upon learning that one asset is the city "Rochetaillee", we would like to send a query to the time-to-flood prediction program.

Were we able to retrieve all of the relevant information, we also assume that a GIS could be used to present the CCMG managers with a graphic such as in Figure 1 as this is the sort of thing that even commercial systems are capable of doing³.

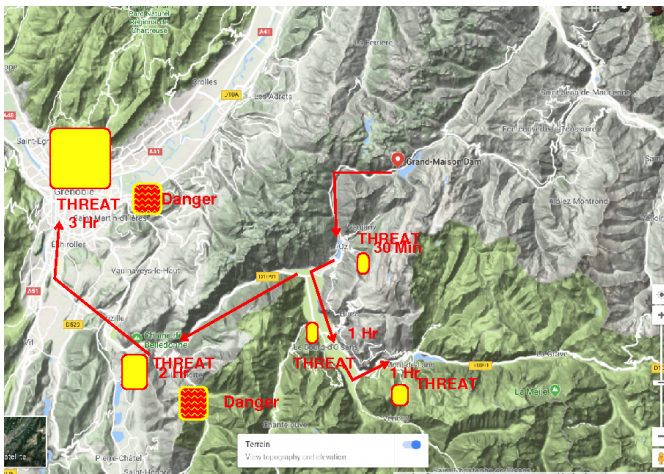


Figure 1: Example Graph for Flood Threat

Assuming these knowledge interchanges can be achieved, especially retrieving the valley names from the GIS, we now move into the response planning domain. The plan returned by a Prefect for a city that would be flooded to a depth of one meter within five hours, might be to evacuate 7000 citizens to another city within that Prefect's Department and another 3000 to some other city in another Department of CCMG's choosing. The means of transport would also be left up to the CCMG, which would then use various analysts to refine each plan it received, determine its efficacy, and refine the sets of plans, removing conflicts. Then CCMG would task various agencies. One task would be perhaps to ask the regional SNCF manager to use trains to transport the 3000 people within the next five hours from the city to be flooded to another in a safe site.

3 MOTIVATING PROBLEM

This is a collage of hypothetical problems considered during the project. Upon being alerted by the CCMG that Rochetaillee and

²For such an alert process, planning *per se* is not necessary: a procedural program could produce such a simple process.

³ERSI: <https://tinyurl.com/y3yl6ury>

other small towns in the valley of La Romanche that they are threatened by the imminent rupture of the Grand Maison Barrage, the Haute-Marne Prefect Mr. Dusserre starts notifying the citizens that they must evacuate and asks the CCMG to refine a plan to move 3000 people by train to Nice.

An industrial chemical plant named IAR in the valley Lac du Verney has several tons of organic peroxides that must be kept refrigerated or explode⁴. After the alert, they request transfer of the chemicals to an unused refrigeration facility in Nice large enough to completely contain all of the hazardous materials and keep them cold. They suggest rail reefer cars could be used.

So now we have two general responses communicated to analysts in the CCMG to refine. Suppose one analyst, working with people in the SNCF develops a plan to re-route three small commuter trains that can carry the evacuees to safety in Nice between, say 13:00 and 16:00 over the main line of rail tracks in the valley. Another analyst working with the SNCF has organized a long train of reefer cars sufficient for moving all of the chemicals and keep them cool for the seven (7) hour trip to the unused refrigeration facility also in Nice. Clearly, the track capacity conflicts should be resolved by the SNCF. But there is more.

First, almost trivially, there is already a big freight train parked near the refrigeration facility in Nice, and this needs to be moved in order for the train with the reefer cars to park there. Less trivially, since the freight train needs to be moved North while the commuter trains are coming South, it needs to be moved before the commuter trains tie up those tracks near Nice.

If this is not noticed, the reefer cars will not be able to reach the refrigeration plant and will explode. If this is noticed, and it should be by the SNCF who should be planning for all track movements, the planner should move the freight train out of the way prior to noon that day. Then the same SNCF planner might think it good to quickly move the reefer car train down there at the same time as the commuter trains are running, as well since the parking space is available by noon. And let us suppose that the SNCF can manage to interleave the trains on the tracks. And secondly, the refrigeration plant in Nice should be activated while the trains are in progress.

But, thirdly, the flooding of the valleys is also going to flood a power plant. The manager has advice that, since it is winter at the height of demand for heating, the power grid will be operating at a lowered capacity. This is yet another response to the imminent disaster. Part of the CCMG's plan is to reduce power to non-critical industry so that the populace can have heating.

The commuter trains run on electricity. There is just enough for them to run. But if the previously unused industrial refrigeration plant is activated, there will not be. Suppose no one has noticed this constraint. Then if the reefer cars are unloaded in Nice prior to the commuter trains finishing their run, either the commuter trains will stop, the chemicals will not be refrigerated and explode, or other people in cities will not have heat. These are bad consequences of an un-noticed conflict among multiple disaster responses.

⁴Such a chemical plant near Houston, TX did explode in 2017 due to loss of power due to flooding from Hurricane Harvey. The plant managers were indicted for not having better contingency plans for loss of refrigeration. <https://www.nytimes.com/2018/08/03/business/arkema-chemical-plant-explosion-texas.html>

The correct plan is to first move the existing freight train from the refrigeration plant; second use the commuter trains to evacuate all of the people, perhaps interleaving the reefer train on the tracks; and then third, only after the evacuation is finished, activate the refrigeration plant and unload the reefer cars. Depending on how long the reefer cars can keep the chemicals cool (which, in Europe, are typically just insulated boxes with no independent refrigeration power), it might be necessary even to load the reefer cars as late as possible. Discovering that this was the necessary sequence after executing plan actions in any other order would be sub-optimal, as previously noted.

This is a simple case, but it would still be easy for various humans in charge of different parts of the plan to miss the electrical power constraints while concentrating on the rail movement logistics.

This example is the first one we have created in this project, and we intend to develop and explore more complex ones as we work with crisis managers. But it is sufficient to suggest the underlying problem that the easiest way to achieve two goals may conflict and a move that is sub-optimal for one of the goals may be necessary in order to provide a plan that is optimal for both.

4 FORMAL PLANNING

As defined in [9] a complete planning technology, given an initial state S_0 , a final state S_f , and a set of all possible actions A_i , each with pre-conditions and effects, can find, if possible, a sequence⁵ of a subset of those actions A_1, \dots, A_n such that S_f is achieved by the effects resulting from execution of A_n and all of the pre-conditions of each action are met as well: a plan is considered valid when all of the pre-conditions for each action are valid in the state when the action is executed.

Planning can be expressed in formal logic and the planner can, using computational logic, produce a plan as a mathematical proof. But there are many ways of achieving this, include generating a plan procedurally and then checking it later for validity. We consider a *planner* any technology that can achieve the goal state with a valid sequence of actions.

It may not be widely appreciated that planning *per se* is non-trivial, especially because of conjunctive goals. The first simple example showing this, called the *Sussman Anomaly*, that of stacking three blocks, was shown in the 70s [12]. A more complex example for supply chain management has been shown to require full planning capability [9]. This example is deliberately complex in order to show how humans might miss the problem and how it would be expensive and time-consuming to identify the problem at execution time. The problem described in section 3 is a planning problem with conflicting conjunctive goals isomorphic to the Sussman Anomaly.

A *planner* able to manage conjunctive goals, typically using either repair or interleaving of actions [1], will produce the correct plan if one is possible.

Here we make no claims about particular technologies. It may be that a linear programming approach, perhaps modified with heuristics, as in [7], could also handle this and other cases of conflicting

conjunctive goals. But it is worth noting that this particular technology, while very general, was targeting evacuation route planning. There are many other responses to many types of disasters that might require a more general form of reasoning.

Our current approach is to apply an interactive planner designed for distributed planning and to coordinate agents performing different parts of the plan. The planner is designed to allow for constraint violations and exceptions that necessitate re-planning. Workflows emerge from the planning dynamically and are revised as needed. This work is based upon the Redux planner [8] and work using the basic Redux system to coordinate multiple agents planning and executing a novel process [5].

Symbolic reasoners, such as Redux [8], are inefficient and typically use too much computational power to solve large problems. However, this particular approach has some advantages. The Redux planner for web services makes it easy to plan to call out to external programs to solve hard problems such as flooding projections and evacuation route planning. It is also easy to incorporate heuristics so as to try certain plans first. This is especially important in disaster planning.

As we noted in our work described in section 2, a dam break disaster will require one kind of process planning whereas a forest fire another. By creating these *meta-plans* in advance, the actual planning for a particular emergency can be greatly optimized.

Also, the Redux approach has another great advantage: it allows planning to proceed even though some constraint violations have been noted. This allows the human planners to reason about how to change the global plan, perhaps at the expense of some individual objectives. This is especially valuable if the situation is simply over-constrained. It may be that either not everyone can be evacuated, or some amount of chemicals will explode, or some people will be without heat for a while. This is typically a political choice that will be presented to some Minister for resolution.

The important point is that by using some form of formal planning in responding to disasters, serious mistakes can be avoided at plan time, rather than corrected at execution time when such correction might have sub-optimal results.

5 GEOSPATIAL PLANNING

In Section 2, we briefly described how a GIS or a spatial database [10, 11] would be used in creating responses to an emergency, based upon meta-plans for doing so. This somewhat blurs the distinction between using a GIS for pre-emergency planning and using a GIS when the emergency is imminent or occurring. Just as in [2], we advocate pre-emergency planning to be part of the contingency plans that the various agencies or the CCMG should draw up prior to an emergency. Both [2, 7] also describe using a GIS for up-to-date information during an emergency: for example, estimating transport capacity.

However, there is little in the literature about how to use a GIS in combination with formal planning to solve such a problem as in section 3. Spatial reasoning is very difficult and the AI community has determined that attacking it with the kind of symbolic reasoning discussed in section 4 is not feasible.

Our approach is to call on a GIS for information as required. A *web service* is a general way to do so, without committing to

⁵Planners need not always produce a sequence: a set of partially ordered actions is often adequate, but they can always be arbitrarily sequenced consistent with the partial ordering.

any specific technology other than the API should have inputs and outputs. Thus, when the alert system needs to know the valleys to be flooded, or the evacuation routes, a GIS can be queried by the planning system. We have demonstrated this in a first prototype. But there are several challenges in using this method in general.

One is simply a consistent representation of time: can the planning system and the GIS sufficiently communicate conditions in some time-frame corresponding to a planning state so that constraint violations can be recognized and repaired? If this can be done, distances should not be an issue as the GIS can be queried as to where a resource traveling at a particular speed will be on a particular path by a certain state of planning.

A plan will have some sequence (actually only partially ordered) $\{S_0, S_1, \dots, S_i, \dots, S_f\}$ states defined by the conditions that are true in that state. A GIS will typically show a set of spatial conditions and boundaries at some time T . So perhaps the flood waters are projected to be, or are sensed to be at some particular depth at some particular time. This should be related to the conditions in some state S_i in which action A_i is planned to be executed. Perhaps A_i is that buses should transport people over some road at this point in planning. One of the pre-conditions of S_i is that this road is accessible.

A simple approach is that if we know the starting time of the various actions and how long they take to execute, we could derive the planned time for the state and then estimate the water depth for that road. We will need to factor in some form of time in any case, so that we know, for example, how long the volatile chemicals could remain in the reefer cars without exploding. Such temporal reasoning has been included in planning previously so relating it GIS query results should be feasible and very sophisticated representations for temporal reasoning in planning have been developed, such as [4]. This can actually be complex and we must find the simplest method that works with GIS data.

Then a variety of reasoning factors would determine if the road is accessible. And during plan execution, if it is sensed that the road too flooded, then a new plan will need to be generated. We don't yet know exactly what queries to a GIS will be necessary in order to generate the kind of information necessary for planning, but we do realize it will have to be something at least similar to GeoSPARQL, with the most promising implementation being Ontotext GraphDB as it provides specific support for ontologies in RDF⁶, such triples being consistent with planners such as Redux.

GIS planning will also play a role in plan repair. Suppose that we find out the water is rising faster than predicted and the road will be impassable when we need it for the buses. We will have to find either a new road that will not be flooded by the time it's needed, or alternative transportation.

Or suppose we determine that the problem in section 3 is over-constrained and we wish to relax the constraint that the train carrying chemicals should not explode. Is there a place along the way that the train can reach that is unpopulated and would be a safe place for the train to explode? What would the GIS query for this look like? In addition, automating generation of such a query in response to an otherwise over-constrained problem has not been

done to our knowledge and would require an advance in planning technology.

Generating the query as above amounts to asking whether some other way of making a prior decision might have been possible, *even though it was not already enumerated as one of the possibilities*. This will call for special heuristics that we have not yet developed though planning has long understood the general problem [3]. Regardless of what planning technology is used, this will be an interesting problem.

And, in order to know whether the road is accessible at certain depth, the system should know what depth would make the road impassable for a bus. In general, we are focusing here on a general issue of *accessibility* of a transportation resource under changing conditions. That is hard enough, but there will be other geospatial features changing continuously that will need to be related to planning states. For example, we will need to know the extent of a particular electric grid and if trains in a certain region would be affected by a reduction in available power. In the abstract, this is also a case of resource accessibility.

Others in the SANCTUM project are addressing the larger semantic issue: a comprehensive formal ontology that will cover all disasters and the various elements of responses to them. This is being done using a semantic network based on [6] by a contractor.

In conclusion, there are hard problems to be solved. Since one of us (Petrie) is a planning expert and one (Voisard) is a geospatial modeling and reasoning expert, we have a good chance of eventual success, especially since we will start with the automation of narrow cases, always advising human managers.

Acknowledgements. This work was carried out in the context of the SANCTUM Project, financed by the French CSFRS.

REFERENCES

- [1] David Chapman. 1987. Planning for conjunctive goals. *Artificial Intelligence* 32, 2 (July 1987), 45. [https://doi.org/10.1016/0004-3702\(87\)90092-0](https://doi.org/10.1016/0004-3702(87)90092-0)
- [2] Thomas J Cova. 1999. GIS in Emergency Management. *Geographical Information Systems* 2 (12 1999), 845–858.
- [3] Vasant Dhar and Casey Quayle. 1985. An approach to Dependency Directed Backtracking using Domain Specific knowledge. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, Vol. 1. Morgan Kaufmann, Los Angeles, 188–190.
- [4] Alfonso Gerevini, Alessandro Saetti, and Ivan Serina. 2005. Integrating planning and temporal reasoning for domains with durations and time windows. In *IJCAI'05 Proceedings of the 19th international joint conference on Artificial intelligence*. Morgan Kaufmann, San Francisco, 1226–1231.
- [5] Sigi Goldman, Andreas Raquet, and Charles Petrie. 1999. Agent-Based Project Management. In *Artificial Intelligence Today (Lecture Notes in AI)*, Vol. 1600. Springer, Berlin, 45:52. https://doi.org/10.1007/3-540-48317-9_14
- [6] Herman Iline. 1986. Un environnement orienté-objets en Prolog. In *Proceedings of CHAM - 86*. Hermes, Marseille, 5.
- [7] Qingsong Lu, Yan Huang, and Shashi Shekhar. 2003. Evacuation planning: a capacity constrained routing approach. In *Proceedings of the International Conference on Intelligence and Security Informatics (LNCS)*, Vol. 2665. Springer, New York, NY, 11. https://doi.org/10.1007/3-540-44853-5_9
- [8] Charles Petrie. 2009. Planning Process Instances with Web Services. In *Proceedings of ICEIS*. Springer, Heidelberg, 31:41. <https://doi.org/10.5220/0002171400310041>
- [9] Charles Petrie. 2016. *Web Service Composition*. Springer, New York, NY. <https://doi.org/10.1007/978-3-319-32833-1>
- [10] Philippe Rigaux, Michel Scholl, and Agnès Voisard. 2002. *Spatial databases - with applications to GIS*. Morgan Kaufmann/Elsevier, Burlington, Massachusetts.
- [11] Jochen H. Schiller and Agnès Voisard. 2004. *Location-Based Services*. Morgan Kaufmann/Elsevier, Burlington, Massachusetts.
- [12] Gerald Jay Sussman. 1975. *A Computer Model of Skill Acquisition*. Elsevier, New York, NY.

⁶<http://www.opengeospatial.org/standards/geosparql>,
<http://graphdb.ontotext.com/documentation/free/>,
 and <https://rdf4j.eclipse.org/about/>