ResQ Freiburg: Deliberative Limitation Of Damage

Alexander Kleiner and Michael Brenner

Institut für Informatik, Universität Freiburg, 79110 Freiburg, Germany

[kleiner, brenner]@informatik.uni-freiburg.de

Abstract. RoboCupRescue is a large-scale multi-agent simulation of urban disasters. In order to save lives and minimize damage, rescue teams must effectively cooperate despite sensing and communication limitations. To accomplish this, ResQ Freiburg uses a twofold strategy: Platoon agents show reactive, yet cooperative behavior, but this can be overridden by deliberative high-level decisions of the center agents. Possible actions are evaluated by prediction modules calculating their instantaneous and long-term effects. Complex sequences of group actions are planned for by a new multi-agent planning method for abstract search spaces that are generated by agent-specific clustering methods.

1 Introduction

The history of search and rescue services has repeatedly shown the difficulty of the coordination of rescue teams in disastrous situations. The RoboCupRescue domain is an attempt to simulate such situations, particularly the aspects of a dynamically changing environment, e.g. the distribution of fire, the bandwidth limitation for communication and the high demand for collaborative sensing and acting in a large scale environment.

The goal is to successfully coordinate teams of agents for the sake of the limitation of damage to people and buildings. There are two classes of agents: platoon and center agents. Platoon agents are small teams that operate in the field, for example fire brigades, and center agents are stations that may instruct the teams in the field with new tasks. Groups of platoon agents take one out of three roles, namely fire brigade, ambulance team and police force. To each kind one center agent is associated.

The domain represents a real multi-agent scenario since most problems are not solvable by a single agent. Fire brigades, for example, rely on blocked roads cleared by police forces in order to reach their target. Buildings are extinguished more efficiently by more than one team, particularly if the fire spreads out in many directions. Moreover, the task is challenging due to limited communication bandwidth, the agents' limited field of view and the difficulty to predict how disasters evolve over time, e.g. how the structure of the city influences the fire spread.

It is the latter feature of the RoboCupRescue domain that most efforts of the 2004 ResQ Freiburg team aim at: In a highly dynamic disaster situation it is a crucial capability of rescue teams to predict future states of the environment as precisely as possible. In contrast to, for example, problems in Classical Planning research, changes do not only result from the agents' own actions in certain areas but also from their *not* acting in others, due to the dynamics of the environment itself. To evaluate the benefit of agents' actions we have integrated new prediction methods in order to facilitate both fast, reactive target selection and deliberative long-term planning of multi-agent behavior. As a result of an extensive randomized testing, data mining and classification (namely the AdaBoost method with Decision Stumps) we have created reliable predictors for the evolution of the health of injured civilians as well as for the influence of extinguishing actions on fires.

Adequately dealing with uncertainty about the *current* state of the world is as fundamental for intelligent behavior as predicting the future. Especially, selecting the current target for an action depends heavily on knowledge about the possible targets' reachability. Using an efficient clustering method, ResQ Freiburg agents can quickly access and evaluate the connectivity of whole regions of a map and base strategic decisions on this information. Furthermore, evidence grids record sensory information, especially about possible locations of injured civilians, thus reducing the time to find and rescue them.

The decision about the execution of actions (or *skills*) is decomposed on the one hand into a reactive part by the platoon agents and on the other hand into a deliberative part by the center agents. The skill selection mechanism of platoon agents uses an evaluation function for possible targets with respect to the current state of the environment. This function is calculated based on the prediction models described above which will be further extended to use hierarchical reinforcement learning [4]. Central agents decide skill execution in the long term by allocating groups of platoon agents to particular tasks. Their decision making is based on a module for state prediction and abstraction that generates the input for a novel multi-agent planner [1].

The remainder of this paper is structured as follows. In Section 2 the inter-agent communication is described in more detail. Section 3 sketches the abilities of platoon and Section 4 the abilities of the center agents, respectively. In Section 5 we summarize our approach and give an outlook to the further development of our team.

2 Communication

To believably model real world disasters the communication of platoon agents in the simulation is limited to the reception and transmission of four messages during each round, respectively. Therefore it is necessary to restrict and compress the amount of broadcasted messages as far as possible. We utilized two methods in order to comply with this requirement. Firstly, each agent memorizes the content of messages perceived during cycles in the past. Therefore it is sufficient only to communicate information about changes. Secondly, the central agents act as message relays between the platoon agents and with other centers.

The first method is realized by an object pool attached to each agent. At start up, the object pool is initialized with the static objects of the map, such as houses and streets. During runtime, newly communicated information about changes in the world, such as fires or blockades, are stored with a timestamp in the data structure and transmitted to the center agent if they were not stored in the pool beforehand. Center agents collect messages transmitted by the platoon agents during each cycle. They filter and merge all collected information into one message that is transmitted to the other center agents and the platoon agents of the same type afterwards.

Furthermore, there exists a special kind of messages for task allocation. Central agents may send an open request for a specific problem to a group of agents. The fire station, for example, might transmit a list of n buildings to a group of m platoon agents situated close to the particular site. Subsequently, the group sends a positive or negative acknowledgment, indicating whether it is possible to perform the task. Central agents may also update an open request e.g. extend the set of buildings that have to be extinguished.

3 Platoon Agents

All platoon agents have the following capabilities:

- World Modeling: Acting efficiently in a chaotic environment like a disaster area is mainly a matter of gathering information and keeping it accurate in spite of high dynamics. The world model each agent maintains is updated not only by its sensory information, but also by the "observation report" compiled by the center agent. Thus agents can share a nearly consistent global world model. Some information may be imprecise: for example, when an agent hears a civilian crying for help the agent can locate the civilian's position only within a range of 30 meters around its current position. Such information is stored in an evidence grid (realized as a hierarchical cluster structure) which allows to reason about probable locations and guides the exploration process.
- Path Planning: A road segment is believed to be either blocked, unblocked or in an unknown state. Depending on what an agent wants to accomplish different versions of an A^*-based path-planning algorithm allow for *safe* planning (use only definitely unblocked roads) or *exploratory* planning (use unblocked or unknown roads).
- Exploration: Agents not having tasks to execute will explore the city to gather information and/or find new opportunities to act. Agents communicate areas chosen for exploration to the others to ensure better dispersion and prevent multiple exploration.
- Target selection: In each cycle platoon agents evaluate possible sites and targets for their actions. The targets are evaluated with respect to their reachability and the expected benefit (cost) of acting (not acting) there.
  The predictive function in use was calculated as a result of an extensive randomized test series that was used to analyze the evolution of properties in the Rescue simulation. We used data mining and machine learning techniques, namely the AdaBoost method with Decision Stumps, to approximate a predictive function both for ambulance agents (prediction of civilian health) and for fire agents (prediction of fire development). In future work, we plan to use the predictive model with hierarchical reinforcement learning [4] in order to learn the optimal target selection function.
  Depending on the current situation and selected target an action (*skill*) is chosen to be performed. Skills are described in a standardized fashion using pre- and postconditions. The precondition is a test that must be fulfilled for the skill to be applicable, the postcondition is another test describing when the skill has successfully been executed. Both conditions correspond to logical descriptions and will be used by a planning component when selecting action sequences for agent teams [1] (cf. Sec. 5). When there is no command from the planning agents or the environment such that a plan does no longer make sense, agents will not stay idle but choose a skill to execute autonomously.
- Acting individually, in team, or on command: all agents have action selection schemes that allow them to chose their actions individually based on their current view of the world. However, those decisions can be overruled when external coordination is necessary: platoons can receive *commands* from their center agent and change their behavior accordingly. For example, fire agents can be commanded to act as a group and extinguish the boundaries of a fire.

3.1 Fire brigades

Fire can only be prevented from spreading by coordinated group action, hence fires are extinguished by teams. The assignment of these teams is carried out by the center agents (cf. Sec. 4).

First, the set of all burning buildings is clustered into separated fire sources. Second, all fire brigades are assigned to extinguish a cluster (selected by the center agent) that is considered the most beneficial cluster to be extinguished. In order to guarantee synchronized actions, one platoon agent is selected among all fire brigades as team leader. The team leader decides from its local view which building in the cluster has to be extinguished next. Ideally, a selected building is burning, is located at the border of the fire area, but not at the border of the map and is likely to be extinguished within a short time. The latter attribute is approximated by a prediction module. The prediction module consists of decision trees that have been trained from a large amount of training cases. The module is capable of answering queries regarding the expected time for extinguishing a building given a team of agents.

After all fires have been extinguished successfully, fire brigades are employed for searching buried civilians in collapsed buildings. Basically, they are performing the same sensing and exploration actions as ambulance agents, as described in the next section.

3.2 Ambulance Teams

Ambulance teams search for injured civilians inside buildings. Besides their own perception they are informed about observations in the field by messages from other ambulance agents and their center agent. All information concerning the whereabouts of civilians is integrated into an evidence grid. The evidence grid represents a probability distribution over the map indicating the likelihood of buried civilians. Multiple hear events in a particular region of the map, for example, increase the probability of civilians being situated in that region; however, other sources of evidences might also be integrated into the grid.

The ambulance teams prefer to dig out civilians with critical injuries, however, not in case a rescue seems to be hopeless. This decision is supported by a prediction module. Given the prediction module, ambulance agents are capable of predicting whether civilians will survive or whether the number of agents performing a rescue mission might not suffice to save a civilian's life. The prediction is based on attributes of the civilian, such as the number of *hitpoints* or the civilian's *buriedness*.

3.3 Police Forces

Police forces are supposed to clear blocked roads to enable other agents to reach their destinations. Thus agents might send an open request for a particular connection on the map to the police station. The police station will delegate the request to one or more of its platoon agents, that will consequently open the requested connection.

Usually it does not suffice to open connections on the map, since connections might still be badly passable due to partially blocked roads. Therefore police forces of the ResQ Freiburg team employ several heuristics to rank roads for clearance, some computed offline, i.e. when the city map is first made public, and some computed during the rescue mission. One offline method predicts importance of road segments by calculating shortest paths between all pairs of nodes on the map. This can be done efficiently with the Floyd-Warshall algorithm [2] (see also [3]). Another offline heuristic increases priorities of roads leading to refuges (important both for civilians seeking shelter and, newly, fire brigades refilling their tanks). Online heuristics include increasing evaluation for roads leading to fire sources and decreasing factors when other police agents are closer to a blocked road.

4 Center Agents

Center Agents are responsible for message routing (see Section 2) and long-term task allocation. Successful task allocation depends strongly on a good prediction of the disaster's extension and the teams' capability to limit the damage. The spreading of fire, for example, depends on the fuel density, e.g. density of houses, the number of floors of each house or the material houses are made of. In order to predict the extent of the fire in the future, one has to take these factors into account. In order to approximate functions for predicting fire spread (as well as other phenomena like evolution of an injured civilian's health) a large set of disaster situations was automatically generated and analyzed with data mining and machine learning methods, namely ensemble learning with the AdaBoost method and decision stumps.

Based on these predictive models, we will use a multi-agent planner to compute complex group plans for teams of agents and thus optimize long-time rewards in the Rescue simulation[1].

5 Summary and outlook

In this paper we introduced the cooperative and deliberative capabilities of the ResQ Freiburg team, in particular cooperative sensing by message routing, prediction-based acting, hierarchical clustering method, and the integration of reactive and deliberative task allocation. At the current stage, deliberation is still very limited in our system. In the future, we want to extend both planning and learning capabilities within the context of hierarchical data structures.

Besides the development of rescue agents, our team plans also to participate in the simulator competition. Precisely, we are looking forward to present the following contributions to the RoboCupRescue simulation community: Firstly a 3D viewer that is capable of visualizing the Rescue system both online and offline. Secondly a novel fire simulator that incorporates a realistic model of fire spread based on the three major ways of heat distribution found in nature[1]. Thirdly a novel collapse simulator that incorporates a 3D physics simulation in order to simulate realistic collapses as well as road blockades as they are typically found after urban disasters. Furthermore we intend to introduce a inter-simulator communication protocol that allows for interactions between simulators, such as the triggering of fires due to a building collapse.

References

1. Michael Brenner. Multiagent planning with partially ordered temporal plans. In *Proc. IJCAI '03*, 2003.

---

[1] A more detailed description of the fire simulator has been submitted to the RoboCup Symposium.

2. T. Cormen, C. Leiserson, and R. Rivest. *Introduction to Algorithms*. MIT Press, 1992.
3. J. Habibi, M. Ahmadi, A. Nouri, M. Sayyadian, and M. Mohammadi. Utilizing different multiagent methods in robocuprescue simulation. Team Description, 2002.
4. A. Kleiner, M. Dietl, and B. Nebel. Towards a life-long learning soccer agent. In *Proc. Int. RoboCup Symposium '02*. Fukuoka, Japan, 2002.