

AtomicOrchid: Human-Agent Collectives to the Rescue (Demonstration)

Sarvapali D. Ramchurn,
Feng Wu,
Nicholas R. Jennings
¹Agents, Interaction, and
Complexity Group
University of Southampton
Southampton, UK
sdr@ecs.soton.ac.uk,
fw16e11@ecs.soton.ac.uk,
nrj@ecs.soton.ac.uk

Wenchao Jiang,
Joel Fischer,
Chris Greenhalgh,
Tom Rodden
²Mixed Reality Lab
University of Nottingham
Nottingham, UK
psxwj@nottingham.ac.uk,
cmg@cs.nott.ac.uk,
jef@cs.nott.ac.uk,
tar@cs.nott.ac.uk

Steve Reece,
Stephen Roberts
³Pattern Recognition Group
University of Oxford
Oxford, UK
reece@ox.ac.uk,
sjrob@ox.ac.uk

1. INTRODUCTION

The coordination of teams of first responders in search and rescue missions is a grand challenge for multi-agent systems research [4]. In such settings, responders with different capabilities (e.g., fire-fighting or life support) have to form teams in order to perform rescue tasks (e.g., extinguishing a fire or providing first aid) to minimise loss of life and costs (e.g., time or money). Thus, responders have to plan their paths to the tasks (as these may be distributed in space) and form specific teams to complete them. These teams, in turn, may need to disband and reform in different configurations to complete new tasks, taking into account the status of the current tasks (e.g., health of victims or building fire) and the environment (e.g., if a fire or radioactive cloud is spreading). Furthermore, uncertainty in the environment (e.g., wind direction or speed) or in the responders' abilities to complete tasks (e.g., some may be tired or get hurt) means that plans are likely to change continually to reflect the prevailing assessment of the situation.

To address these challenges, a number of algorithms and mechanisms have been developed to form teams and allocate tasks. For example, [6] and [1], devised centralised and decentralised optimisation algorithms respectively to allocate rescue tasks efficiently to teams of first responders with different capabilities. However, none of these approaches considered the inherent uncertainty in the environment or in the first responders' abilities. Crucially, to date, while all of these algorithms have been shown to perform well in simulations (representing responders as computational entities), none of them have been *exercised* to guide *real* human responders in real-time rescue missions. Thus, it is still unclear whether these algorithms will cope with real-world uncertainties (e.g., communication breakdowns or changes in wind direction), be acceptable to humans (i.e., be clear for humans and take into account their capabilities), and actually augment, rather than hinder, human performance.

In this demo we present AtomicOrchid, a novel game to evaluate team coordination under uncertainty using the concept of mixed-reality games. AtomicOrchid allows an agent, using a task planning algorithm, to coordinate, in real-time, human players using mobile phone-based messaging, to complete rescue tasks efficiently. By so doing, it provides a platform to study the notions of flexible autonomy and agile teaming for human-agent collectives.

Appears in: Alessio Lomuscio, Paul Scerri, Ana Bazzan, and Michael Huhns (eds.), *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, May 5-9, 2014, Paris, France.

Copyright © 2014, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

2. THE DISASTER SCENARIO

We consider a disaster scenario in which a satellite, powered by radioactive fuel, has crashed in a sub-urban area.¹ Debris is strewn around a large area, damaging buildings and causing accidents and injuring civilians. Moreover, radioactive particles discharged from the debris are gradually spreading over the area, threatening to contaminate food reserves and people. Hence, emergency services (including medics, soldiers, and fire-fighters) are deployed to evacuate the casualties and key assets (e.g., food reserves, medication, fuel), each requiring different teams of responders, before they are engulfed by the radioactive cloud. In what follows, we describe the use of a planning agent at headquarters to help coordinate the team.

2.1 Human-Agent Collaboration

In line with practice in many countries, we assume that the first responders are coordinated from a headquarters (HQ) headed by a human coordinator H . In our case, H may be assisted by an autonomous task planning agent PA (more details in Section 2.2), that can receive input from, and direct, the first responders. Both H and PA can communicate their instructions (task plans to pick up targets) directly to the responders using an instant messaging system (or walkie talkie). While these instructions may be in natural language for H , PA instructs them with simple requests such as "Pick up target X at position Y with team-mates Z " messages. In turn, the responders may not want to do some tasks (for reasons outlined above) and may therefore simply accept or reject the received instruction from PA or H .² However, H can query the responders' decisions and request more information about their status (e.g., fatigue or health) and goals (e.g., meeting with team-mate at position X or going for task Y). Instead, if a task is rejected by the responders, PA records this as a constraint on its task allocation procedure and returns a new plan. Thus on the one hand, richer interactions are possible between H and the first responders than between them and PA . On the other hand, PA runs a sophisticated task allocation algorithm that can compute an efficient allocation, possibly better than the one computable by H (particularly when many responders need to be managed).

In our demo, the focus will not be on the algorithm run by PA , but focus on the interactions between human players and the planner agent. Hence, in the next subsection we provide a brief overview

¹Given the invisibility of radiation, it is possible to create a believable and challenging environment for the responders to solve in our mixed-reality game (see Section 3).

²While some agencies may be trained to obey orders (e.g., military or fire-fighting), others (e.g., transport providers or medics) may not be trained to do so.

of the task allocation/path planning algorithm used by PA and then go on to describe the AtomicOrchid game.

2.2 Planning Agent Algorithm

Previous agent-based models for team coordination in disaster response typically assume deterministic task executions and environments [6, 7]. However, in order to evaluate agent-guided coordination in a real-world environment, it is important to consider uncertainties due to player behaviours and the environment (as discussed in the previous section). Given this, we develop a new representation for the task allocation problem in disaster response that does take into account such uncertainties. More specifically, we represent this problem using a Multi-Agent Markov Decision Process (MMDP) that captures the uncertainties of the radioactive cloud and the responders' behaviours. We model the spreading of the radioactive cloud as a random process over the disaster space and allow the actions requested from the responders to fail (because they decline to go to a task) or incur delays (because they are too slow) during the rescue process. Thus in the MMDP model, we represent task executions as stochastic processes of state transitions, while the uncertainties of the radioactive cloud and the responders' behaviours can be easily captured with transition probabilities.

At each decision step, we assume PA fully observes the state of the environment by collecting sensor readings of the radioactive cloud and GPS locations of the responders. Given a policy of the MMDP, a joint action can be selected and broadcast to the responders. We next describe the game within which PA uses its algorithms to interact with human players.

3. THE AtomicOrchid GAME

We adopt a serious mixed-reality games approach to counteract the limitations of computational simulations [3]. The impact of emotional and physical responses likely in a disaster, such as stress, fear, exertion or panic remains understudied in approaches that rely purely on computational simulation [2]. In contrast, our approach creates a realistic setting in the sense that participants experience physical exertion and stress through bodily activity and time pressure, mirroring aspects of a real disaster setting [5]. This, in turn, provides greater confidence in the efficacy of behavioural observations regarding team coordination supported by a planning agent.

3.1 Interacting with the Planning Agent

To permit the interaction between First Responders and PA , the former are equipped with a 'mobile responder tool' providing sensing and awareness capabilities in three tabs (geiger counter, map, messaging and tasks; see Figure 1). The first tab shows a reading of radioactivity, player health level (based on exposure), and a GPS-enabled map of the game area to locate fellow responders, the targets to be rescued and the drop off zones for the targets. The second tab provides a broadcast messaging interface to communicate with fellow first responders and the commander H . The third tab shows the team and task allocation dynamically provided by the agent PA that can be accepted or rejected. Notifications are used to alert both to new messages and task allocations.

3.2 Running the Demo

For the purposes of a demo at AAMAS, it was envisaged that conference attendees will be recruited as players of the game during the demo session and that the game will be run live. As it stands, the venue for the conference is in a location where there will be poor GPS reception and lack of space for a live game to take place. We will therefore endeavour to run sessions of the game in the Garden of the Observatory (7 mins walk from the conference centre) provided approval is granted for this. Otherwise, we will run videos of the demo and allow attendees to play with the mobile app and interact with the planning agent in simulated scenarios.



Figure 1: Mobile first responder and HQ interfaces.

For live runs, the headquarters will be located in the conference centre. H will have at her disposal an 'HQ dashboard' that provides an overview of the game area, including real-time player location information (see Figure 1). The dashboard provides a broadcast messaging widget, and a player status widget so that the responders' exposure and health levels can be monitored. H can further monitor the current team and task allocations to individual responders by PA (by clicking on a button). Crucially, only H and PA have a view of the radioactive cloud, depicted as a heatmap ('Hotter' (red) zones correspond to higher radiation levels).

First responders will be assigned a specific type: medic, firefighter, soldier, and transporter. Their mission is to evacuate all four types of targets: victim (requires medic and fire-fighter), animal (requires medic and transporter), fuel (requires soldier and fire-fighter), or other resource (requires soldier and transporter). The first responders are supported by (at least) one person (H) in a centrally located HQ room, and the planning agent PA that sends the next task (as described earlier) to the team of responders.

The game lasts a maximum 30 mins. Participants will be briefed and asked to consent to participate. They will be presented with a demographic questionnaire to record gender, occupation, experience of using smartphones and level of map navigation skills. At the end of the briefing in which mission objectives and rules will be outlined, responder roles will be randomly assigned to all participants (fire-fighter, medic, transporter, soldier).

Acknowledgements

This work was carried out as part of the ORCHID project funded by EPSRC (EP/I011587/1).

4. REFERENCES

- [1] A. Chapman, R. A. Micillo, R. Kota, and N. R. Jennings. Decentralised dynamic task allocation: A practical game-theoretic approach. In *Proc. of AAMAS 2009*, pages 915–922, May 2009.
- [2] J. Drury, C. Cocking, and S. Reicher. Everyone for themselves? a comparative study of crowd solidarity among emergency survivors. *The British journal of social psychology*, 48(3):487–506, 2009.
- [3] J. E. Fischer, M. Flintham, D. Price, J. Goulding, N. Pantidi, and T. Rodden. Serious mixed reality games. In *Mixed Reality games Workshop at the CSCW 2012*, 2012.
- [4] H. Kitano and S. Tadokoro. Robocup rescue: A grand challenge for multiagent and intelligent systems. *AI Magazine*, 22(1):39–52, 2001.
- [5] Pan American Health Organization. Stress management in disasters, 2001.
- [6] S. D. Ramchurn, A. Farinelli, K. S. Macarthur, and N. R. Jennings. Decentralized coordination in robocup rescue. *The Computer Journal*, 53(9):1447–1461, 2010.
- [7] P. Scerri, A. Farinelli, S. Okamoto, and M. Tambe. Allocating tasks in extreme teams. In *Proc. of AAMAS*, pages 727–734. ACM, 2005.