

ARES 2: A Tool for Evaluating Cooperative and Competitive Multi-agent Systems

Jörg Denzinger and Jordan Kidney

Dep. of Computer Science, University of Calgary
{denzinge, kidney}@cpsc.ucalgary.ca

Abstract. The Agent Rescue Emergency Simulator (ARES) system provides a simplified rescue scenario similar to Robocup Rescue for use in the educational or research fields when evaluating multi-agent systems. The environment within ARES abstracts ideas down to key features while still allowing for a wide range of scenarios requiring different skills to be presented to the agents. This is achieved by combining different features together and changing the configuration of the environmental rules incorporated in ARES. With the simplified environment ARES can be used for quick evaluations of concepts based upon the results from a range of different configurations of the system. The results from these tests can be used as a basis for further experiments when transitioning the work from a theoretical level to more real world scenarios, where the requirement for a more complex system is needed. Our newest version of ARES, ARES 2, allows for scenarios that range from cooperation between all agents to strong competitiveness of agents or even agent teams.

1 Introduction

The Agent Rescue Emergency Simulator (ARES) system (see [4, 2]) provides a simplified rescue scenario similar to Robocup Rescue (see [5, 6]). The Robocup initiatives helped to expand the field of multi-agent systems, bringing many different areas together in a single competition. From this initiative, a set of simulators was produced for the competitions, allowing the agents to interact without robotic bodies. These simulators also provided a consistent environment for evaluating teams entered in the competition. Slowly these simulators have started to find their way into the research and educational fields (see [1, 3]), but several problems have surfaced. One such problem is the fact that the simulators were created for a very specific application, mainly the environment needed for the competition. As a result, the range of concepts that can be demonstrated is limited to what the competition required. What is needed now is to extend upon the success of the simulators from the Robocup initiatives, but with the main goal of providing a flexible testbed for the research and educational communities, allowing for the quick evaluation of new concepts in multi-agent systems. Our solution to this problem is the ARES system, which follows the same ideas as Robocup rescue, but focuses on providing a very flexible environment that can be tailored to particular research or educational needs.

The basic ideas realized in ARES 2 are as follows. The environment simulated within ARES 2 is a city after an earthquake has struck (similar to Robocup rescue).

The components of the environment have been simplified down to what we consider to be the basic features relevant for multi-agent systems. As such we do not concentrate upon providing an environment to simulate real world scenarios, but one that concentrates on providing an area for demonstrating a range of different concepts. This allows us to break some real world rules with the main goal of allowing a range of concepts that can be demonstrated through the use of the system.

Next for the agents there are two basic tasks, finding survivors and rescuing them, that emphasize the need for independent exploration, coordination, and cooperation between agents. Agents have a rather limited set of actions, including the ability to communicate and they act within a world that is more or less known to them. The agents acting in a scenario in ARES 2 might be one team or they can be in several teams, thus allowing various combinations of cooperation aspects and competition aspects.

A key concept in ARES 2 are the so-called world rules that result in configuring key features of the ARES 2 world scenarios. Selecting a set of world rules essentially creates basic requirements on the agents and the scenarios using this set of rules. And if we assume that we want to develop agents and their interactions in such a way as to optimize their task performance, then different sets of world rules create different agent goals and different requirements on what is a good behavior. This achieves a high flexibility for the use of ARES 2 to evaluate teams of agents, since we can target particular aspects of multi-agent systems by choosing the right set of world rules. Features that are influenced by world rules include competitiveness between agent teams, cost of communication, precision of world information and observability of the world, or fulfilling the basic survival needs of agents.

So far, ARES and ARES 2 have been mostly used in educational settings, although several researchers have expressed an interest in them for evaluating agent teams using particular concepts. In the educational setting, the possibility of changing world rules has proved very useful, since plagiarism essentially is eliminated. With the addition of world rules influencing the competitiveness and interaction possibilities of several agent teams acting in the same scenario in ARES 2, concepts like trust and communication outside of the own team can now also be evaluated with ARES simulations.

2 ARES 2: Basic Ideas

For a more technical explanation of the ARES System, please refer to [4, 2].

2.1 The Environment

The environment simulated within ARES 2 is based upon a city that has been hit by an earthquake. This is very similar to Robocup Rescue, but the objects within the world and the actions that can be taken have been simplified down to what we consider to be the basic features relevant for multi-agent systems. Now depending upon the configuration of the system the agents will either have the ability to work with all the agents currently connected to the system or only with agents that have been declared as part of their group. This allows for the creation of environmental rules that can push the interaction between teams of agents from a cooperative scenario to a more competitive environment. Agents within the system interact within a two dimensional grid environ-

ment built up of squares. Each square in the grid represents a single point where an agent can influence the environment directly. The agents "jump" from square to square as they move throughout the world. Each square in the world has the following properties:

1. It can be of one of the following types: (1) Normal - agents can safely move onto the square. (2) Fire - the entire square is on fire. (3) Killer - represents a zone where agents will die (like holes, etc ...) and finally (4) Charging - a zone that can be used to regain lost energy
2. It consists of a stack of layers (build up of material), where each layer holds only a single object. Currently there are only three types of objects in the ARES system: (1) Rubble - represents material that has to be removed by the agents. Each rubble object has as associated value the number of agents that are needed to remove this object. This forces agents to coordinate and come together upon the same square in the world at the same time to remove the object. (2) Survivor - a single person that can be saved in the world. When they are saved, the survivors are simply "beamed" to safety. (3) Survivor Group - this is the same as a survivor, it just allows for multiple survivors to be located in a single layer on the stack.
3. Each square also has an associated move cost value; this value indicates the cost in energy for moving onto the square from any direction

In general, the interaction of an agent with ARES 2 consists of the agent sending messages indicating its actions (including any communication actions to other agents) to ARES 2 when they are informed that it is their turn to run. Once ARES 2 has given all agents a chance to run and send an action it then applies all the actions to the current state of the environment. Next each agent is informed about the results of their actions (if a reply is expected). This process is repeated until the specified number of simulation steps has been met. For the agents in the environment, each action that they take has a corresponding consequence based upon the type of action that is being taken. In the ARES 2 system, agent actions fall into two groups: (1) Actions that influence the environment directly. These actions include things such as moving around, removing rubble, etc. With the execution of these actions the agents lose energy, as a result the agents must manage their energy usage and when they must recharge. (2) Communication/Observational actions. Actions such as communication and observing squares in other areas of the world fall into this group. With the execution of these actions the agents have to deal with things such as communication delays and distortion of information.

2.2 World Rules, Environment Features, and Their Connection to Multi-agent Systems

As already stated, a major goal in the development of ARES and ARES 2 was to provide a lot of possibilities of how to configure the environment to allow for different foci on the general problems of multi-agent systems. ARES 2 achieves this by implementing different so-called world rules that describe how effects in the environment can be or have to be achieved by the agents. This is possible in ARES 2, because of the rather simplified worlds that agents are acting in within ARES 2.

By combining the selected world rules for different world features, the resulting environment will require rather different strategies by the agents to be successful. And these strategies will focus on different aspects of the agents, of their implementation, and of how the agents interact. The world features new in ARES 2 (compared to ARES) are how to deal with different agent teams interacting in the same simulation run. For example, by allowing multiple distinct teams to participate in the simulation the agents are presented with a competitive environment to work in. By changing the way scoring is done for saving survivors the ways that agents from different teams work together will change. By giving the point to all who participate it gives the opportunity for all agents to work together and gain a mutual benefit. If we just change the scoring to give the point to the team who had the most participants in the save action then the environment becomes a little bit more hostile for agents from different teams to work together and always gain a mutual benefit. The main world features influenced by world rules in ARES 2 are as follows:

1. Score for saving survivors - changes to this feature effect the competitiveness of the environment.
2. Single or multiple teams in the simulation - cooperative vs competitive environment.
3. Maximum number of agents required to remove rubble - causes the agents to deal with resource allocation issues by having to get other agents to help in the removal of the rubble object.
4. Energy Control - agents need to deal with monitoring their energy level and regaining energy back as it is needed. In some configurations of the system this becomes a resource allocation issue as the agents may have to move onto a specific square in the world to regain energy.
5. Communication - changes in how the agents communicate effect everything from agent models and cooperation concepts to resource allocation. For example, by limiting the amount of communication an agent can do in one step, the number of steps needed to get help for a rubble removal can become higher and the general task more difficult.

It should be noted that individual features can be influenced by several world rules, which allows us to provide a large variety of feature instantiations.

3 Experimental Evaluation

So far, our experiences with ARES and ARES 2 have been in the educational field. The ARES system has been used for the past three years at the graduate level –once at the undergraduate level– for courses on multi agent systems at the University of Calgary. Feedback from the students has shown a positive image of the ARES system. The system has provided a consistent way to compare the agents produced by different teams in the class. Each time the class has been taught different configurations of the world rules have been used to present the students with a different challenge while still working in the common rescue environment. Another reason for the change in the world rules has been to prevent the "reuse" issue of students using everything from

agents produced in previous versions of the class. The thought behind these variations is to make an environment where agents that in one year would not work properly in the newly changed environment. This would mean that they would be slower at finishing tasks, or would be at a greatly elevated risk of dying off faster. The students could look at old designs but would have to develop their own ideas and implementation to deal with the new environment. Each time the multi-agent course has been taught there has been a noticeable improvement in the quality of the agents produced by the students. Students in the latest version of the class were able to take into account faults and observed problems of the teams from previous years and could improve upon the concepts successful in previous years in their own way. Overall many positive results have come from using the ARES and ARES 2 system in the course and not only were the students rather enthusiastic about applying what they have learned in the lectures to a scenario where even the instructor did not know exactly what the best concepts would be, concepts have been developed that have the potential to be used in more realistic testbeds or scenarios.

4 Conclusion

We presented the ARES 2 system, a testbed for evaluating multi-agent systems that was motivated by the Robocup initiatives, but aims at providing much more flexibility in the environments it offers to the agents. ARES 2 in particular introduced the ability of having scenarios with several competing teams interacting at the same time and allowing for some cooperation between these teams. While this is rather common in the real world, the known testbeds for multi-agent systems, including the Robocup testbeds do not offer this possibility. Our use of ARES and ARES 2 as testbed for systems build as assignments for multi-agent courses showed that the simplified worlds of ARES make it rather easy to build agents for ARES. In the future, we want to make ARES even more flexible by adding more world rules, respectively more instantiations of the current ones. We also hope to attract other researchers to use ARES not only in educational settings, but also to evaluate their research results at a higher level.

References

1. Paul Buhler, José M. Vidal. Biter: A platform for the teaching and research of multiagent systems' design using robocup, Proceedings of the Robocup International Symposium, 2001.
2. Jörg Denzinger, Jordan Kidney, Melissa Bergen. Teaching Cooperation in Multi-Agent Systems with the help of the ARES System, Proc. WCCCE-03, Courtenay, 2003.
3. F. Heintz, J. Kummeneje. Simulated robocup in university undergraduate education, technical report., Department of Computer and Information Science, Linköping University, 2000.
4. Jordan Kidney. ARES Website <http://www.cpsc.ucalgary.ca/kidney/ARES> (as viewed on March 8, 2004).
5. RoboCup Rescue. <http://jelly.cs.kobe-u.ac.jp/robocup-rescue/> (as viewed on Nov 24, 2003).
6. RoboCup. <http://www.robocup.org/> (as viewed on Nov 24, 2003).