The ALADDIN Project: Intelligent Agents for Disaster Management

N.M. Adams[1] M. Field[2], E. Gelenbe[3], D.J. Hand[1], N.R. Jennings[4], D.S. Leslie[5],
D. Nicholson[6], S.D. Ramchurn[4], S.J. Roberts[7], A. Rogers[4]

1. Institute for Mathematical Sciences, Imperial College London, South Kensington Campus, London SW7 2PG, UK
2. BAE Systems, Insyte, Lyon Way, Frimley, Camberley, Surrey GU16 7EX, UK
3. Dept. Electrical and Electronic Engineering, Imperial College London, South Kensington Campus, London SW7 2AZ, UK
4. Electronics and Computer Science, University of Southampton, Southampton SO17 IBJ, UK
5. Dept. Mathematics, University of Bristol, University Walk, Bristol BS8 1TW, UK
6. BAE Systems, Advanced Technology Centre, Filton, Bristol BS34 7QW, UK
7. Dept. Engineering Science, University of Oxford, Parks Road, Oxford OX1 3PJ, UK

ABSTRACT

ALADDIN [1] is a multi-disciplinary project that is developing novel techniques, architectures, and mechanisms for multi-agent systems in uncertain and dynamic environments. The application focus of the project is disaster management. Research within a number of themes is being pursued and this is considering different aspects of the interaction between autonomous agents and the decentralised system architectures that support those interactions. The aim of the research is to contribute to building more robust multi-agent systems for future applications in disaster management and other similar domains.

1. INTRODUCTION

The disaster management domain is characterised by: distributed control; uncertainty, ambiguity, imprecision and bias; multiple stakeholders with different aims and objectives; and limited resources which continually vary. In the ALADDIN project, a disaster management system is viewed as being composed of autonomous, reactive and proactive agents that can sense, act and interact in order to achieve individual and collective aims. These agents need to be able to make the best use of available information, be flexible and agile in their decision making, cognisant of the fact that there are other agents, and adaptive to their changing environment. This requires a multi-disciplinary approach, in particular:

- Filtering and data fusion methods for estimating relevant state variables, such as the position of rescue vehicles and wounded civilians
- Decision-making and machine learning methods for determining actions in response to states, such as when and where to route a particular rescue vehicle
- Multi-agent systems, game theory, and mechanism design methods to manage the interaction between multiple actors and to model collective behaviour
- System architecture studies of different agent organisations and information exchange topologies, for example centralised, hierarchical or decentralised systems

The challenge is to integrate these elements in order to develop decentralised data and information systems that can operate effectively in highly uncertain and dynamic environments. This is not just a research challenge; it is a key requirement for many industrial and commercial organisations, exemplified by the disaster management application domain.

This paper summarises the four research themes that make up the ALADDIN project – Applications, Decentralised System Architectures, Individual Agents, and Multiple Agents.

2. APPLICATIONS

The Applications research theme is focusing on the disaster management domain. This domain was chosen because it is important in its own right, it poses demands from a functional and integrated systems point of view, and it requires timely decision making and actions in highly uncertain and dynamic situations. To ensure the specific methods and techniques developed across the ALADDIN research programme fit together to form a coherent whole, the project has developed three main software demonstrations:

- **Situation Awareness [2]**. This demonstration makes use of a network of weather sensors located on the South Coast of England. These sensors measure a range of environmental variables and publish up-to-date sensor measurements on separate web pages. Weather sensors have immediate application in the disaster management context. They exhibit challenging correlations and delays whose physical processes are not well understood, and they are subject to network outages that generate instances of missing sensor readings. To facilitate the autonomous collection of sensor data by information processing algorithms, each sensor web page was supplemented with machine-readable RDF data. Further, in order to visualise the sensor data and the results of the information processing algorithms, a Java software implementation was developed. This is a prototype of the software that could run on the mobile computer or PDA carried by first responders in order to provide situation awareness support at the scene of a disaster. A display screen from the demonstration is shown in Fig.1.

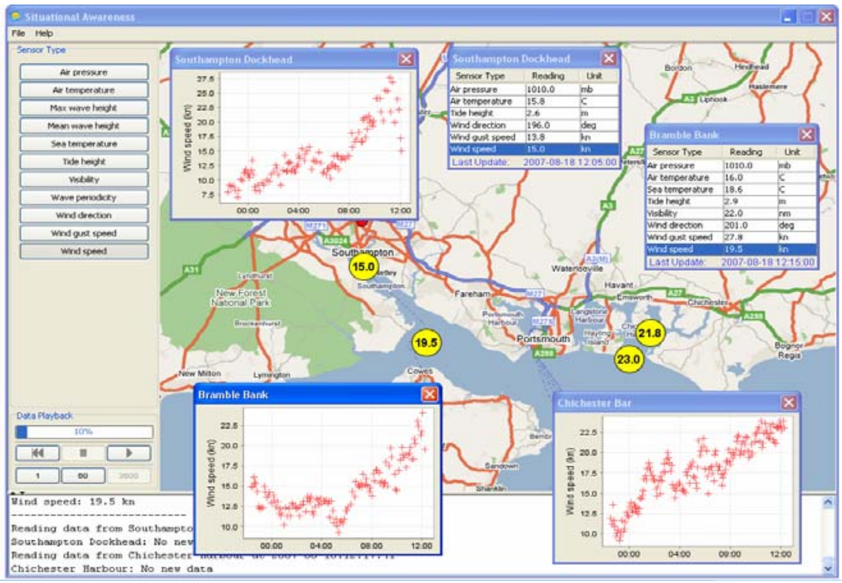

Fig. 1 Situation Awareness demonstration display

- **Building Evacuation [3]**. This demonstration considers emergency response in a tall building, where each floor looks similar, there are few escape routes, and many choke points. The building is modelled in simulation as a graph where the nodes represent physical locations and the edges represent paths between those locations (edge weights describe path lengths and other properties). Civilians and rescuers (agents) within the building move around by ‘hopping’ between the nodes. A virtual fire is created in the building and a real wireless sensor network is integrated in the simulation to monitor the spread of the fire. LEDs, controlled by an external fire emulator, provide input to the sensors about the intensity of the fire. The system will be used to demonstrate a number of intelligent agent-based evacuation strategies as well as the effect of different network architectures and techniques on an emergency response situation. Some display screens from the demonstration are shown in Fig.2.

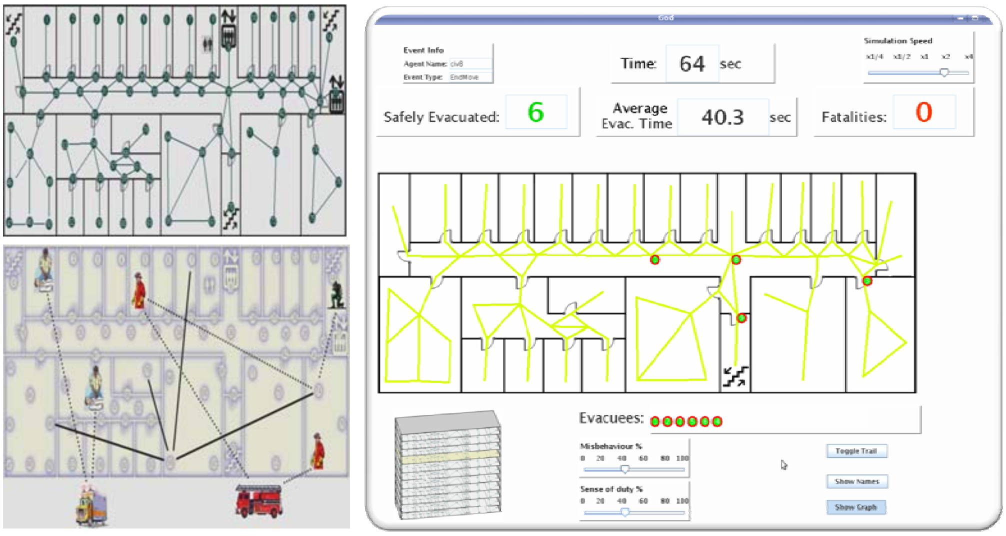

Fig. 2 Building evacuation demonstration display

- **Robocup Rescue [4]**. This demonstration builds on an existing simulation tool developed by the Robocup Rescue project following the 1995 Kobe earthquake. The tool simulates civilian behaviour, traffic, fires, building collapse and blockades, following the aftermath of an earthquake. Police, ambulance and fire service agents have to rescue civilians, clear blockades, and extinguish fires. The ALADDIN project has built extra functionality into Robocup Rescue to increase its flexibility. This includes a function to degrade sensing capability (with noise or faults) to test the robustness of various rescue strategies and a function to declare agents as active (in moving vehicle) or inactive (in damaged stationary vehicle) to test whether a strategy can deal with dynamic situations where a team of agents has to reconfigure to manage a disaster. This extended Robocup Rescue system will enable ALADDIN to demonstrate novel methods for challenging problems such as: path planning in the presence of collapsed buildings and road blocks; efficient communication in constrained bandwidth networks; multi-source information fusion in real time, and effective coordination within and between teams. Display screens are shown in Fig.3.

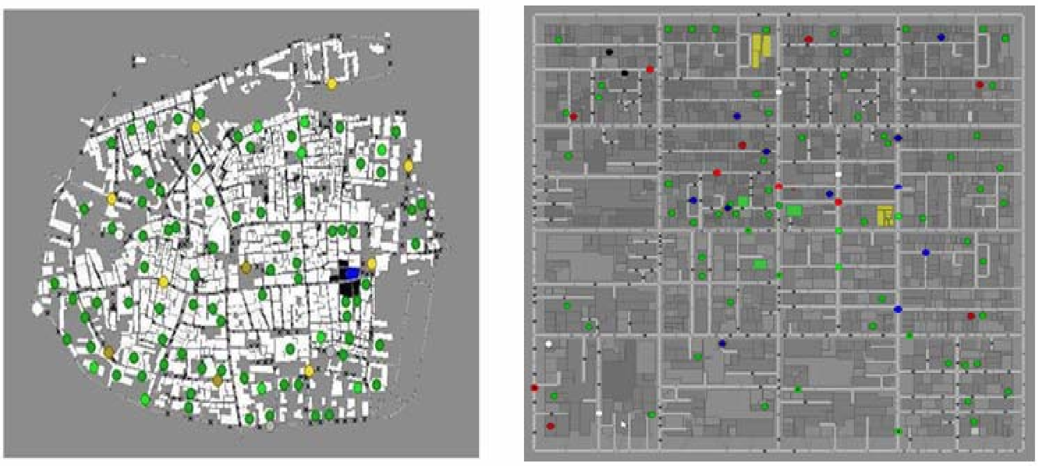

Fig. 3 Robocup Rescue demonstration display

3. DECENTRALISED SYSTEM ARCHITECTURES

The disaster management domain is complex because it is composed of multiple functional levels and multiple organisational structures. This research theme is concerned with defining some candidate architectural options and (with a particular focus on decentralised system architectures) evaluating their performance in different disaster management scenarios.

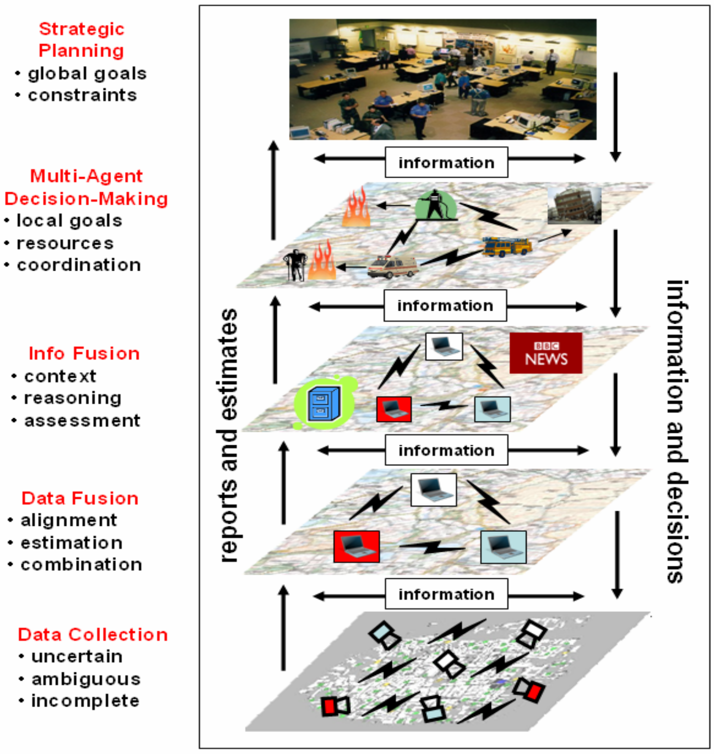

Fig. 4 Main functional levels

The main functional levels are shown in Fig.4. At the lowest level there is a data collection function where multiple sensors acquire data about the domain, such as the position of rescue vehicles and civilians. At the next level, a data fusion function associates and combines this data to reduce its uncertainty. The next level reasons about the data in the disaster management context to assess the situation and its threats. This forms the basis at the next level for agents (i.e., rescuers) to decide how they should act to create an effective coordinated response. Finally, at the top-level, there is a strategic planning function that sets goals and priorities which drive this response.

Broadly speaking, data in the form of sensor reports and fused estimates propagate up the processing chain, while information and decisions propagate back down. This is a closed loop process because the agents typically decide how they should act in order to collect further data to process in support of their response objectives.
A number of alternative organisational structures are possible in the disaster management domain. Some of these are shown in Fig. 5. The conventional organisation type is centralised. Data is collected by a number of distributed sensors but then transformed into information, knowledge, and actions, at one or more dedicated central processing stations. This organisational set-up is unsuitable for a large-scale disaster management system because

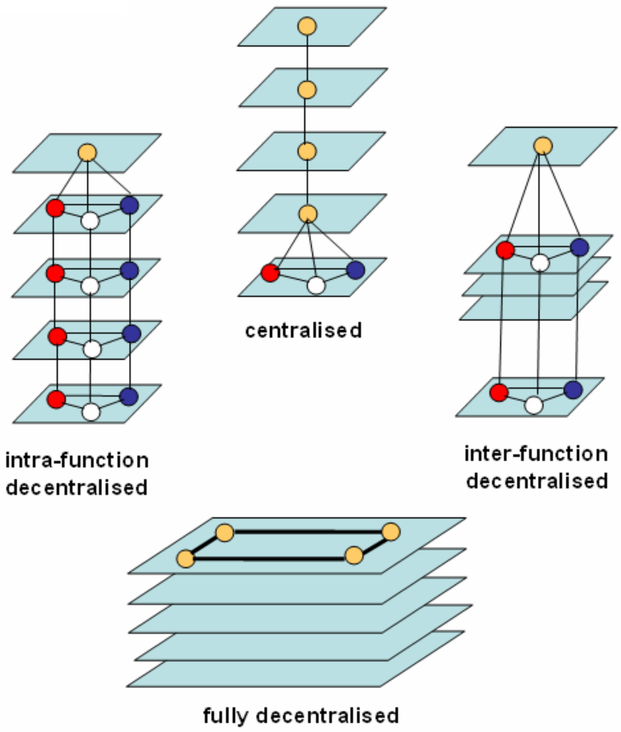

Fig. 5 Organisational structures

the strict vertical flow of information between the processing stations and the actors would create bottlenecks in the system and reduce response times. Consequently, the ALADDIN project is focusing its research around decentralised system architectures. The fully decentralised architecture in Fig. 5 is decentralised in two respects:

- the functional levels are conflated to create intelligent agents that can sense, think and act for themselves, without relying on dedicated central processing stations
- the intelligent agents communicate with each other in an opportunistic manner, peer-to-peer, instead of through a central broadcast communication station

Full decentralisation offers the potential advantage of robustness and scalability, leading to shorter response times, at the expense of a more complex information management problem due to the open and fluid nature of the system.

Of course, there are many architectural options in between the fully centralised and fully decentralised extremes. Fig. 5 illustrates two of them. In an intra-functional decentralised architecture there may be separate functional levels in the disaster management organisation, but fully decentralised processing and communication at each level. For example, a decentralised sensor network in the disaster zone for tracking rescue vehicles and civilians, and a decentralised decision-making network outside the disaster zone involving the police, fire and ambulance command cells engaged in communication and cooperation. In an inter-functional decentralised architecture, there may be some conflation of the functional levels to empower the police, fire and ambulance agents with more autonomy to decide and act, but with some limitations on cross-service communication to simplify information management.

The initial work in this theme, which set out to determine the range of issues and variables that govern possible disaster management architectures and how these options can be compared and contrasted, is now complete. The work has entered a demonstration phase, using the extended Robocup Rescue environment developed in the Applications theme, to evaluate their relative merits across a representative range of disaster management scenarios.

4. INDIVIDUAL AGENTS

This research theme is concerned with developing the theory, tools, and techniques for individual agents, which form the building blocks of a distributed data and information system. This is a significant challenge for two main reasons. First, it needs to take a holistic view of the individual agent. Thus, each individual agent must:

- fuse information obtained from its environment in order to form a coherent view of its world that is consistent with other agents
- make inferences over this world view to predict future events
- plan and act on its conclusions in order to achieve its objectives given these predictions

These activities need to be performed on a continuous basis because of the many feedback loops that exist between them. Second, each agent must operate in this closed loop fashion in an environment where there is a significant degree of uncertainty, resource variability, and dynamism, and there are multiple other agents operating under a decentralised control regime. To be effective in such contexts, a number of important research challenges need to be addressed. In particular, uncertainty, ambiguity, and incompleteness will be present at all levels within the system, which calls for a robust and flexible approach throughout.

The ALADDIN project has made significant progress in a number of key areas, for example:

- **Delayed Data [5].** In a distributed dynamic sensing environment, due to bandwidth restrictions, it is usual for sensor reports to be received some time after they were generated. Such delays can be complete (the full observation vector is delayed), partial (parts of the observation vector are delayed), continuous, or random. Estimation techniques for dealing with delayed data generally fuse all of the data without considering the processing/improvement trade. Research in ALADDIN has considered signal-to-noise ratio and application needs as metrics for deciding *when* and *which* delayed measurements to fuse. In practice, an adaptive fusion mechanism is motivated by potential savings in computational resources and power consumption.

- **Missing Data [6].** In a disaster management system it is particularly likely that some of the sensing or communication components will fail due to the hostile operating conditions. This can give rise to intermittent gaps in certain data streams and compromise the performance of individual agents. In ALADDIN this problem has been approached using the powerful and flexible technique of Gaussian Processes for prediction. A Gaussian Process enables Bayesian inference about functions – in this case functions over time. The algorithmic details are described in the reference. A representative result is shown here (Fig. 6) to demonstrate the utility of this technique.

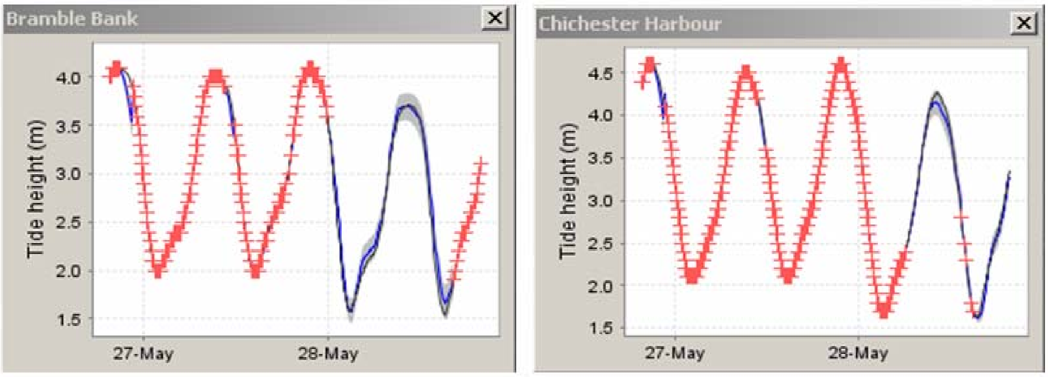

Fig. 6 Gaussian Process prediction and regression of real weather data

The Situation Awareness application demonstrator, described earlier, provides an excellent real-world testing ground for many of the techniques developed in the Individual Agents theme, including Gaussian Processes. Fig. 6 displays a sample of sensor observations of tide height, acquired by two separate weather stations over a period of about 36 hours. These are indicated by the red crosses. Several episodes of missing data from both sensors are evident, caused by network outages. The Gaussian Process predictions (mean and error bar) are also shown (blue line and grey shading, respectively). These accurate predictions exploit temporal correlation in the separate sensor observations, as well as cross-correlation between the sensors observations.

- **Spurious data [7].** Individual agents must also be able to handle spurious data from sensors which are malfunctioning or faulty in some unknown way, otherwise this data can dominate and generate erroneous world views and misguided actions. In order to deal with the problem, the ALADDIN project has developed Generalised Covariance Union (GCU). This novel approach exploits probabilistic bounds on the reliability or trustworthiness of multiple data sources to merge their data in a statistically consistent manner. In doing so, GCU unifies the well known techniques of multiple hypothesis data reduction and Kalman filtering under a theory of trust. The GCU solution also provides a statistically sound framework for learning the trustworthiness of sensors.

The Robocup Rescue system has been used in ALADDIN to test the ability of GCU to consistently track fire spread within a sensor network in the presence of sensor faults. Fig. 7(a) displays the true temperature evolution over time (black line), and the estimated temperature (mean – blue line, error – grey shading) using a naïve approach that assumes all data is reliable. Clearly, these estimates are inconsistent. Fig. 7(b) displays the GCU result, which is consistent.

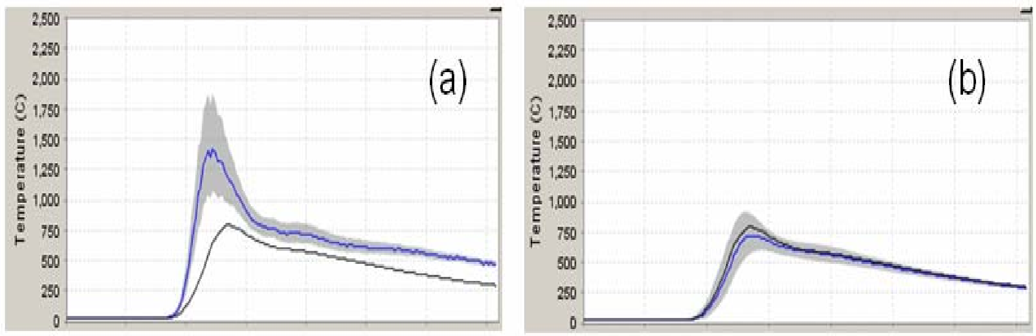

Fig. 7 Estimation of fire spread (a) inconsistent, (b) consistent (GCU-based)

- **Decision-Making.** Research is also underway in ALADDIN to link inference and prediction to decision-making and control. The formulation and performance of techniques such as Reinforcement Learning and Multi-Arm Bandits will be explored, given the dynamic, uncertain nature of the disaster management domain, and the requirement for individual agents to operate effectively under decentralised control.

5. MULTIPLE AGENTS

This research theme is primarily concerned with the way in which the various autonomous agents within the system interact with one another in order to achieve their individual and collective aims. It covers three main types of activity:

- how the interactions of the autonomous agents can be structured such that the overall system exhibits certain sorts of desirable properties
- the class of methods that such agents can use to coordinate their problem solving when the system is operational
- how the interactions of such agents can be modelled and simulated in order to determine the macroscopic behaviour of the overall system based on the microscopic behaviour of the participants

To tackle the above activities, a number of techniques are being used to analyse and shape the interactions between multiple agents in order to achieve the overall system wide properties. These include (but are not limited to):

- **Auction Methods [8]**. Auctions are an efficient means of dynamically allocating limited resources in the presence of multiple stakeholders with minimal communication requirement. In the disaster management domain the different stakeholders correspond to the police, fire and ambulance services, rescuers, and civilians. The resources include the service vehicles and equipment, routes, sensors, bandwidth, etc. Research in ALADDIN is focusing on the design of auction strategies (e.g. bidding and payment strategies) and mechanisms to achieve fair, rational, and efficient system goals. It has examined dynamic auctions and distributed auctions which are modelled as multiple auctions closing simultaneously. The extended

Robocup Rescue system is being used to demonstrate the research since it is a dynamic environment with multiple stakeholders. Fig.8 displays a screenshot.

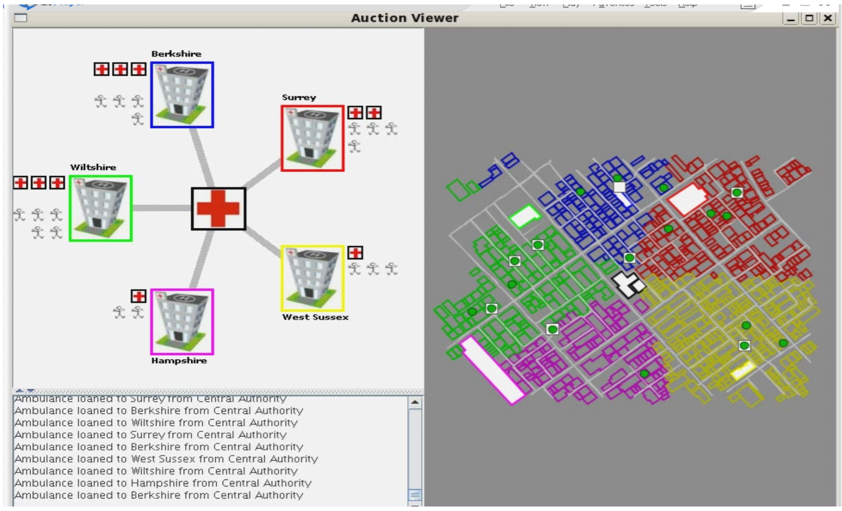

Fig. 8 Robocup Rescue demonstration of Auction methods

- **Coalition Methods [9]**. Coalitions are formed when sub-sets of agents are assigned to a specific task. For example, in disaster management, an ambulance and fire vehicle could form a coalition to rescue a civilian trapped in a burning building. In a large-scale disaster management system, the challenge is to search through the enormous space of potential coalition structures quickly and with guarantees on solution quality. Research within the ALADDIN project has developed such an algorithm, and it offers several orders-of-magnitude improvement on current state of the art. The algorithm has been applied in the Robocup Rescue system to an ambulance allocation problem.
- **Learning in Games [10]**. The problem of multiple autonomous agents trying to learn how to coordinate their actions to meet an objective can be cast as a game. In the ALADDIN project the requirement is to solve such games in a decentralised fashion. A number of approaches are being explored with a view to understanding their theoretical relationships and ultimately their scope and limitations in the disaster management application context. The main approaches include Fictitious Play, Coordinated Reinforcement Learning, and Distributed Constraint Optimisation.

6. CONCLUSIONS

The ability of emergency responders to deal with disasters is being severely stretched given the challenges posed by increasingly complex and non-traditional operations. The need is for decentralised command and control systems that are more flexible and can respond more quickly to new information. This applies at the strategic, operational and tactical levels. However, decentralisation in a dynamic environment requires new and innovative technical solutions that can provide timely and robust inference and decision support. The ALADDIN project is researching such solutions and a number of them have been described here. These are characterised by their multi-disciplinary nature, bringing together tools and techniques from statistics, data fusion, machine learning, and game theory. The intention is to contribute to building more robust multi-agent systems for future applications in disaster management as well as other related domains (e.g. military command and control in urban operations).

Acknowledgements

The ALADDIN (Autonomous Learning Agents for Decentralised Data and Information Networks) project is jointly funded through a BAE Systems and EPSRC (Engineering and Physical Sciences Research Council) strategic partnership.

References

1. Autonomous Learning Agents for Decentralised Data and Information Networks, (www.aladdinproject.org)
2. A. Rogers, M.A. Osbourne, S.D. Ramchurn, S. Roberts, and N.R. Jennings. "Information agents for pervasive sensor networks." In the 4th International Conference on Sensor Networks and Systems (PerSeNS), 2008
3. A. Filippoupolitis, L. Hey, G. Loukas, S. Timotheou, and E. Gelenbe. "Augmented reality simulation of emergency response in a multi-storey building." *ACM Symposium of Applied Computing (SAC 08)*, Track on Self-Organisation in Pervasive Distributed Systems (SOPDS 08), 2008
4. www.robocuprescue.org
5. D.K. Tasoulis, N.M. Adams, and D.J. Hand. "Should delayed measurements always be incorporated in filtering?" Proc. 15th IEEE Conference on Digital Signal Processing, 264-267, 2007
6. M.A. Osbourne, S.J. Roberts, A. Rogers, S.D. Ramchurn and N.R. Jennings. "Towards real-time information processing of sensor network data using computationally efficient multi-output Gaussian processes." *Information Processing in Sensor Networks (IPSN 08) Conference*, 2008.
7. S. Reece and S. Roberts, "Generalised Covariance Union with an application in time-delayed measurement fusion" *IEEE Trans. on Aerospace and Electronic Systems*, 2008
8. E.H. Gerding, R.K. Dash, D.C.K. Yuen, and N.R. Jennings, "Optimal bidding strategies for simultaneous Vickrey auctions with perfect substitutes" Proc. 8th Intl. Workshop on Game Theoretic and Decision Theoretic Agents, Hakodate, Japan, 10-17, 2008
9. T. Rahwan, S. D. Ramchurn, V.D. Dang, A. Giovannucci, and N.R. Jennings, "Anytime optimal coalition structure generation" Proc. 22nd Conf. pf Artificial Intelligence (AAAI), Vancouver, Canada, 1184-1190, 2007
10. D. S. Leslie, "Game-theoretic learning for sensor networks" Autonomous *Agents and Multi-Agent Systems Conference (AAMAS 07)*, In 1st International Workshop on Agent Technology for Sensor Networks (ATSN-07), 2007