

ALADDIN: Research into Multi-Agent Solutions for Decentralised Information Systems

Simon Case (BAE Systems, Advanced Technology Centre)

Nicholas R. Jennings (University of Southampton)

Robert Johnston (BAE Systems, Integrated System Technologies)

Simon.case@baesystems.com nrj@ecs.soton.ac.uk robert.johnston2@baesystems.com

ABSTRACT

ALADDIN (Autonomous Learning Agents for Distributed and Decentralised Information Networks – <http://www.aladdinproject.org/>) is a five year partnership between industry and academia to develop methods and techniques for decentralised information systems. Currently in its fourth year, its aim is to use autonomous software agents to address challenges of situational awareness, inference and co-ordination in decentralised systems. To this end, this paper outlines the primary goals of ALADDIN and reviews the work undertaken up to the current time in the context of information exploitation.

INTRODUCTION

The case for decentralisation of military Command and Information Systems has been made persuasively in a sequence of monographs published by the Command and Control Research Program of the DoD (the Office of the Assistant Secretary of Defense (NII)) The monograph of most relevance here is Understanding Command and Control by David S. Alberts and Richard E. Hayes, 2006 [1]. Related thinking contributed to BAE Systems' formulation of the need for decentralised technologies in complex information systems [5] and defined the requirements for the ALADDIN project.

It is clear that such decentralisation offers robustness, in the sense of having no single point of failure, in support for graceful degradation and also in the ability to maintain and even enhance performance as sub-systems are added. However a stronger need has been argued for by Alberts and Hayes [1]. This is that in systems of non-trivial complexity a centralised actor will not have access to sufficient information to make local decisions that are effective, let alone optimal, across the system under consideration. The local actors have information that is critical to decision making and action in their immediate environment. Technologies to enable effective information exploitation and management in complex command and information systems will need to be decentralised and the technology described here enables the collaboration of actors (agents') to achieve common, system level goals while maintaining autonomy in decision making and action. The multi-agent system techniques that are being generated through the ALADDIN project allow best use of sensor, and other data sources, and effective delivery of information as required by the stakeholders throughout the system.

ALADDIN (Autonomous Learning Agents for Distributed and Decentralised Information Networks – <http://www.aladdinproject.org/>) is a five year partnership between industry and academia to develop methods and techniques for decentralised information systems. Currently in its fourth year, its aim is to use autonomous software agents to address challenges of situational awareness, inference and co-ordination in decentralised systems. To this end, this paper outlines the primary goals of ALADDIN and reviews the work undertaken up to the current time.

Report Documentation Page				Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.						
1. REPORT DATE OCT 2009		2. REPORT TYPE N/A		3. DATES COVERED -		
4. TITLE AND SUBTITLE ALADDIN: Research into Multi-Agent Solutions for Decentralised Information Systems				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) BAE Systems, Advanced Technology Centre				8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)		
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited						
13. SUPPLEMENTARY NOTES See also ADB381582. RTO-MP-IST-087 Information Management - Exploitation (Gestion et exploitation des informations). Proceedings of RTO Information Systems Technology Panel (IST) Symposium held in Stockholm, Sweden on 19-20 October 2009., The original document contains color images.						
14. ABSTRACT ALADDIN (Autonomous Learning Agents for Distributed and Decentralised Information Networks http://www.aladdinproject.org/) is a five year partnership between industry and academia to develop methods and techniques for decentralised information systems. Currently in its fourth year, its aim is to use autonomous software agents to address challenges of situational awareness, inference and co-ordination in decentralised systems. To this end, this paper outlines the primary goals of ALADDIN and reviews the work undertaken up to the current time in the context of information exploitation.						
15. SUBJECT TERMS						
16. SECURITY CLASSIFICATION OF:				17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 6	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified				

ALADDIN: Research into Multi-Agent Solutions for Decentralised Information Systems

In [5] it was noted that designing systems in the Network Enabled Capability (NEC) context raises a number of challenges to current state of the art technologies in terms of dealing with:

- **Structure:** The analysis process is open-ended and gives rise to issues such as opportunism, planned perception and planned resolution of ambiguity. The outputs are diverse and range from identification of specific threats to trends and relationships.
- **Uncertainty:** It is not always possible to be precise at the outset about what information is required and there are a wide range of information sources (qualitative as well as quantitative) with different levels of uncertainty, and associated confidence.
- **Dynamism:** Time is a key factor. This requires us to balance accuracy against speed where there is a wide range of information sources (qualitative as well as quantitative) with different time constants.

In particular, NEC requires information technologies that are able to operate when many events are occurring at once; the situation is evolving on many time-scales and the environment is non-cooperative. Against this background it is necessary to co-ordinate different assets under incomplete information and the decision-making and control need to be distributed amongst those assets. This requires information fusion to be distributed and take place at many levels. Furthermore, solutions to these problems need to be flexible, scalable, robust, traceable and compatible with existing architectures and systems. ALADDIN was formed to undertake fundamental research on these challenges.

ALADDIN RESEARCH

To address the aforementioned challenges, ALADDIN has investigated a large number of techniques and algorithms from both a single-agent perspective and from the perspective of large sets of multiple interacting agents. In the first instance the issue is how an individual agent responds to the inputs from the world around it. To this end, ALADDIN has applied statistical techniques to the selection, interpretation and fusion of those inputs in order to determine an appropriate action, and the onward prediction of future states based on its current picture of the world. In the latter case, the issue is one of how such agents can act in an environment where there are many other agents and the individual and joint goals of these agents require them to cooperate, coordinate and negotiate.

Single agents

The project has made significant progress in dealing with delayed, missing and spurious data. In the context of a single agent attempting to best exploit the available information, it can be useful to know when fusing input data will improve the current picture or whether it can simply be ignored. The work on delayed measurements [14] describes a set of experiments for understanding how the arrival of out-of-sequence data affects the performance of a variety Kalman-filter based data fusion methods. The results show that although fusing all measurements provides optimal results, this comes at the price of greater processing and storage demands. On the other hand performance degrades gracefully as we reduce the proportion of delayed values fused. Depending on the method and other factors such as the signal to noise ratio, the results show that 80% performance can be achieved with only 50% of the delayed measurements.

One means by which information obesity can be addressed is by reducing the amount of information which needs to be observed. [2] describes a method for selecting which variables need to be monitored in a set of incoming data. The method has been developed for streaming data and does not assume that the correct variable set is static but addresses the *Adaptive Variable Selection of Regression* problem in which the most suitable set of variables is dynamic over time.

The method is based on the *Lasso algorithm* [12] which is an approach for linear regression known to be appropriate for sparse data environments such as those with large input vectors. The lasso algorithm requires the mean and co-variance of the covariates to be known and to compute these the new method uses the well-known EM algorithm [4] applied to the most recent observations. The results show that almost equivalent powers of prediction can be achieved when using only 10% of the data.

Effective exploitation of information can also depend on the user's ability to extrapolate forward from current data. [5] describes an online principled Bayesian method for multi-variate regression which is used both to fill in gaps between data points, and also to enable short term prediction to support decision making. The approach uses the technique of Gaussian Processes [9] in which a multivariate Gaussian prior distribution is assigned over the outputs of the regression problem. Analytic posterior distributions for outputs of interest which are conditional on whatever data has been collected are then produced. Importantly, these posterior distributions are also Gaussian, and thus we have a predictive mean, and also a variance that explicitly represents uncertainty. Users of this technique therefore not only know what the predicted values will be at a given time step, but also how good a prediction this is.

The implementation of this technique within the Aladdin project has enhanced the approach by developing an efficient method which is able to handle online data, and which can adaptively update the model as new observations are made. Because of this online approach the method can also be used for active data selection. Thus when prediction uncertainties grow beyond a certain point, the approach can suggest which data input needs to be queried to best reduce uncertainty.

The method has been validated against weather data from a set of weather sensors in the South of England. The results have been impressive with the technique learning correlations between data inputs and providing good prediction performance. Given the principled nature of the approach, it is expected that this method will provide similar performance on other data sets, being of particular use when data from a variety of sources is correlated.

Multiple agents

This theme addresses the challenge of how large scale agent systems can be designed so that aspects of overall system behaviour can be predicted even though the individual components may be autonomous agents designed and implemented by different stakeholders. Hence we have developed systems using techniques with guaranteed performances across relevant metrics. From an exploitation point of view these are important results as it provides an approach to managing emergent system properties. The algorithms have been drawn from a number of disciplines including game theory, market-based approaches, statistical methods and reinforcement learning techniques.

Some important results have included new algorithms for creating coalitions of agents in time scales several orders of magnitude smaller than the current state of the art [7]. Decentralised control of autonomous nodes has also been addressed by a number of means. In particular, factor graph type methods for control have been investigated and [13] extends the scalable max-sum algorithm to enable operation with continuous control parameters. A final stream of the research is investigating game-theoretic methods and [3] describes an application of these techniques to decentralised dynamic task allocation.

From an information exploitation point of view, a critical issue in a decentralised approach to complex information systems is how to value the information from other parts of the system. [10] describes an approach to combining information in sensor environments in which some of the faults may be modelled but some may be unanticipated. This approach will be described in more detail in [6] which will also be presented at this conference, so is not described here.

ALADDIN: Research into Multi-Agent Solutions for Decentralised Information Systems

One important aspect of the value of a piece of information is the level of trust attached to it. Indeed, within a decentralised system in which the provider of information is not known to the user of that information, one way of enhancing its exploitation is to have the ability to assign a level of trust to the information. In this way users can exploit information which they may have otherwise ignored, and equally may ignore information which otherwise would have required processing. [11] describes a new technique for assessing and transmitting the trustworthiness of nodes in a network which has application to many different types of contract including one in which information is being supplied. Computational trust approaches should ideally enable a means of expressing the trustworthiness of agents which (i) allows agents to update their estimates of a supplier as they acquire direct experience (ii) provides a natural framework for agents to express their uncertainty about this trustworthiness and (iii) allows agents to exchange, combine and filter reputation reports received from other agents.

The approach taken models the trust of an agent as a vector of estimates corresponding to the probability that any dimension of the contract will be successfully fulfilled and exploits a covariance matrix which describes the uncertainty and correlations in these probabilities. These probabilities are calculated from the direct experience of the agent using a formalism based on the Dirichlet distribution. This representation addresses points (i) and (ii) above. To communicate this information around the system and address point (iii) the method uses the sufficient statistics of the Dirichlet distribution. Specifically for each pair of dimensions the agents communicate a vector of contract outcomes, each of which is the sufficient statistic for the corresponding Dirichlet distribution. Each agent is therefore able to send out information about its own trust estimates in messages of size $2d(d-1)$ where d is the number of contract dimensions. Other agents can then incorporate these messages with their own direct experience to gain a more precise estimate of trustworthiness.

Of course, such an approach can lead to the problem of *data incest* in which because the provenance of each individual contract outcome becomes lost in the aggregation process, observations of contract outcomes can be included more than once. This can lead to bias and over confidence. To avoid this, the approach uses the concept of *private* and *shared* information. The agent decomposes its contract outcome vector into two parts. The private part is information which it knows it has not shared with any other agent, and the shared information which it has either transmitted or received from other agents. By following a procedure depending on whether the agent is sending or receiving the message as described in the paper, it is guaranteed that contract outcomes in the shared vector are never combined with those of another agent's shared vector, and hence outcomes which originate from the same agent are never combined together.

EVALUATION AND DEMONSTRATION

In order to provide a means for benchmarking the algorithms against each other and existing techniques, three demonstrator frameworks have been developed, each emphasising different aspects of NEC.

First, the situational awareness demonstrator is aimed at assessing data and information fusion techniques. It fuses the inputs from a set of independent weather sensors which are located in the Solent. This data has the characteristics required for assessment in that it is a set of distributed sensors with correlated, but not equal, inputs and importantly we have a known ground truth. In addition, further data can be taken from a wider set of sensors some of which are known to provide unreliable data, which allows us to exercise our methods for assessing the reliability of sensor information. A web accessible version of the demonstrator can be found at <http://situation.ecs.soton.ac.uk/>. Both the Gaussian Process [7] and the variable selection work [2] have been tested and demonstrated on this data.

Second, to demonstrate and assess the methods corresponding to multi-agent systems, we have used and extended the "Robocup Rescue" environment (<http://www.robocuprescue.org/>). In this

demonstration a set of blue-light emergency vehicles act collaboratively to minimise the loss of life and damage in the aftermath of an earthquake. The different types of entities (police, fire and ambulance) are heterogeneous in that they are capable of different actions and hence can affect the overall outcome differently. Furthermore, in this demonstrator there is no central control and so each entity must act as best it can with only a limited view of the situation given by its own sensors and communications with near neighbours. This demonstrator therefore allows us to exercise our research in the areas of multi-agent behaviour. This has been a very fruitful testing and demonstration environment and demonstrations have included the work on data fusion with untrustworthy sensors [11], the application of auction techniques for resource allocation, and the construction of teams of blue-light agents using the coalition formation work [8].

The final demonstrator is the building evacuation simulator. This shows the spread of a fire through the floor of a building that contains software agents whose aim is to assist both the occupants who are leaving the building and the responders who need to enter the building. In particular, the purpose of this is to exercise decentralised control algorithms at a smaller scale than RoboCup rescue and to extend the mix of actors to specifically include the actions of human participants. An example of this is shown in [15] where a heuristic approach is taken to maximising connectivity between people in a building by decentralised control of a set of autonomous wireless nodes.

CONCLUSIONS

ALADDIN was formed with the intent of providing fundamental research to address the technology needs inherent in a network enabled information system. The major challenge that is being addressed is the discovery of techniques that allow command and information systems to be effective in decentralised systems. The move to decentralised systems is essential if information management and exploitation is not to breakdown [1].

To date, a multi-disciplinary team has made significant progress in addressing a number of the fundamental challenges through the creation of a toolbox of algorithms and methods. These algorithms have been implemented and evaluated in a number of challenging demonstrator frameworks and their potential in such domains has been clearly shown. From an information management perspective, the algorithms have addressed several different aspects of information exploitation. These have included a technique to know when to include delayed data. Another method referred to as *adaptive variable selection for regression* enables prediction with only 10% of the original data and is adaptive to changes in the data stream. Gaussian processes are a means of exploiting data by principled prediction based on what is currently known, its use has enabled development of online active data selection techniques to query input data to best reduce uncertainty. In a decentralised environment in which information providers may not be personally known to information users, a key issue is how to represent *trust* in the system, and the Aladdin team have developed a method for expressing and circulating trust measures around a system, which avoids data incest issues.

These algorithms and others developed with Aladdin have subsequently been applied to the areas of:

Sensor Integration - the management and control of sensors as information sources in Command and Information Systems

Threat evaluation and weapon assignment - collaborative resource assignment without a centralised decision maker

Defence Logistics - exploiting information at all point of the logistics chain to increase the efficiency of the system.

REFERENCES

- [1] David S. Alberts and Richard E. Hayes (2006) "Understanding Command and Control" CCRP . http://www.dodccrp.org/files/Alberts_UC2.pdf
- [2] C. Anagnostopoulos, N. M. Adams and D. J. Hand (2008) "Deciding what to observe next: adaptive variable selection for regression in multivariate data streams" *Proc. ACM Symposium on Applied Computing, Cear Brazil*.
- [3] A. Chapman, R. A. Micillo, R. Kota and N. R. Jennings (2009) "Decentralised Dynamic Task Allocation: A Practical Game-Theoretic Approach" *Proc. 8th International Conference on Autonomous Agents and Multiagent Systems (AAMAS '09)*, Budapest, Hungary, 915-922.
- [4] A. Dempster, N. Laird and D. Rubin. Maximum Likelihood from Incomplete Data via the EM Algorithm *Journal of the Royal Statistical Society Series B (Methodological)* 39(1):1-38, 1977
- [5] N. R. Jennings, R. Johnston and A. Wright (2006) "Network Enabled Capability: Decentralised Coordination of Autonomous Agents to Achieve Operational Goals" *NATO RTO-MP-IST-055 Information Fusion for Command Support*, The Netherlands. (<http://www.rta.nato.int/Pubs/RDP.asp?RDP=RTO-MP-IST-055>)
- [6] D. Nicholson, Defence Applications of Agent-based Information Fusion Technology. IST-087 Symposium on Information Management/Exploitation
- [7] M. A. Osborne, S. J. Roberts, A. Rogers, S. D. Ramchurn and N. R. Jennings (2008) "Towards real-time information processing of sensor network data using computationally efficient multi-output Gaussian Processes" *Proc. Int. Conf. on Information Processing in Sensor Networks*, St. Louis, USA.
- [8] T. Rahwan, S. D. Ramchurn, V. D. Dang, A. Giovannucci, and N. R. Jennings (2007) "Anytime optimal coalition structure generation" *Proc. 22nd Conf. on Artificial Intelligence*, Vancouver, Canada, 1184-1190
- [9] C.E. Rasmussen and C.K.I. Williams, *Gaussian Processes for Machine Learning*, MIT Press 2006
- [10] S. Reece, S. Roberts, C. Claxton and D. Nicholson "Multi-sensor fault recovery in the presence of known and unknown fault types" 12th International Conference on Information Fusion, Seattle, USA, 2009
- [11] S. Reece, A. Rogers, S. Roberts and N. R. Jennings (2007) "Rumours and reputation: evaluating multi-dimensional trust within a decentralised reputation system" *Proc. 6th Int. J. Conference on Autonomous Agents and Multi-agent Systems*, Hawaii, USA.
- [12] R. Tibshirani. *Regression Shrinkage and Selection via the Lasso*. *Journal of the Royal Statistical Society Series B (Methodological)* 58(1): 267-288, 1996
- [13] R. Stranders, A. Farinelli, A. Rogers and N. R. Jennings (2009) "Decentralised coordination of mobile sensors using the max-sum algorithm" *Proc 21st Int. Joint Conf on AI*, Pasadena, USA.
- [14] D. K. Tasoulis, N. M. Adams, and D. J. Hand (2007) "Should Delayed Measurements Always Be Incorporated In Filtering?" *Proc. 15th IEEE Conference on Digital Signal Processing*.
- [15] S. Timotheou and G. Loukas (2009) "Autonomous networked robots for the establishment of wireless communication in uncertain emergency response scenarios" *Proc. 24th ACM Symposium on Applied Computing, Track on Intelligent Robotic Systems*, Honolulu, Hawaii.