Network Enabled Capability: Decentralised Coordination of Autonomous Agents to Achieve Operational Goals

Prof. Nick Jennings
Electronics & Computer Science
University of Southampton
Southampton SO17 1BJ
UK

Dr. Robert Johnston
BAE Systems
Integrated System Technologies
Lyon Way
Frimley, GU16 7EX
UK

Prof. Andy Wright
BAE Systems ATC
Sowerby Building
PO Box 5
Bristol BS34 7WQ
UK

nrj@ecs.soton.ac.uk

robert.johnston2@baesystems.com

andy-src.wright@baesystems.com

ABSTRACT

The move from platform centric planning to network enabled capability poses not only organisational challenges but also challenges to the technology that will be needed to enable networked operations. Within this context, this paper outlines the central role that autonomous software agents, that coordinate their activities in flexible ways, have to play in providing robust system solutions. The major research challenges are addressed in order to take this promising technology forward and to make it suitable for NEC application.

The interactions of the various autonomous agents within the system, that are necessary in order to achieve their individual and collective aims, can be analysed and designed using techniques from Game Theory and Mechanism Design. Game theory is exploited because it has developed powerful tools for analysing decision making in decentralised open systems with multiple autonomous agents. Recently, these tools have been tailored to computational settings to provide a principled foundation for building multiple agent systems. This tailoring gives rise to the field of computational mechanism design.

A whole system approach to the requirements and the assessment of the technologies has been adopted. Of prime important for driving this research programme are:

- ***Information Assurance:** Currently in networked systems, security and trust are key issues with no robust solution. Therefore assurance of information over dynamic networks requires tackling dynamic and cross coalition, multi-level security systems.*
- ***Self-management of Sensor Networks:** The complexity of future systems means that a degree of sub-system autonomy will be required for self-management of the networks and its assets.*

1.0 INTRODUCTION

It is widely recognised that there is an on-going revolution in military systems, which emphasises need to continue and reduce critical time-delays in command and control but in wider and more varied operational contexts. The manipulation, management and distribution of information play an essential role in effecting this flexible capability. In particular, the strong drive for the provision and exploitation of a Network Enabled Capability (NEC) is one manifestation of this change. To this end, the 2003 UK Defence White Paper (presented to Parliament by The Secretary of State for Defence in December 2003) [1] stated: “our focus is now on delivering flexible forces able to configure to generate the right capability in a less predictable and more complex operational environment. This will require us to move away from simplistic

Jennings, N.; Johnston, R.; Wright, A. (2006) Network Enabled Capability: Decentralised Coordination of Autonomous Agents to Achieve Operational Goals. In *Information Fusion for Command Support* (pp. 15-1 – 15-12). Meeting Proceedings RTO-MP-IST-055, Paper 15. Neuilly-sur-Seine, France: RTO. Available from: http://www.rto.nato.int/abstracts.asp.

platform-centric planning, to a fully "networked enabled capability" able to flexibly exploit effects-based planning and operations".

Against this background, this paper outlines the central role that autonomous software agents, that coordinate their activities in flexible ways, have to play in this endeavour. Specifically, we highlight the major research challenges that need to be addressed in order to take this promising technology forward and to make it suitable for NEC applications. Before doing this, however, we outline the major requirements and stages of NEC as they apply in our context.

1.1 NEC Research Requirements

NEC is built upon the principle that military operation will be enhanced if there is an environment in which:

- Information is shared
- All users have consistent information on the battle space
- Decisions are made collaboratively
- Effects in the battle space are synchronised

To achieve this, the following areas must be addressed:

- **Three dimensions of NEC:** There is need to cover the NEC dimensions of - People, Information and Networks. Research must focus on technologies and infrastructures for connecting people to the relevant information through the existing and future networks and equipment.

- **MoD Lines of Development** (LoD): Provide an important test for completeness of the solutions that are enabled by a research or development programme. Thus there should be coverage of the mapping on to Lines of Development, which should be considered as threads running through programmes not as specialist stovepipes.

- **Delivery of Relevant Information:** The three Dimensions of NEC and the LoD emphasise the delivery of relevant information to the appropriate personnel. This is central to the concept of transformation through interconnection. In particular ensuring that future systems can adopt to different cultural and coalition environments

- **A Whole System Approach:** Related to the above, but along a different axis, is the need to pursue research programmes that look at the whole system. Technologies and solutions must be applicable at the Tactical, Operational and Strategic level and apply to the movement between levels. Of specific important here are:

    - **Information Assurance**: Currently in networked systems, security and trust are key issues with no robust solution. Therefore assurance of information over dynamic networks requires tackling dynamic and cross coalition, multi-level security systems.

    - **Self-management of Sensor Networks:** The complexity of future systems means that a degree of sub-system autonomy will be required for self-management of the networks and its assets. The emphasis will not be on the sensors themselves but on their network readiness.

1.2 NEC Research Specification

To ensure the best use of available data, information and knowledge, there is a need to take a total system view of the data and information fusion process and to ensure that the system architecture solutions support this view. Specifically, this must include the feedback between sensing, decision making and action processes, which within the military world is called the OODA loop. In a data and information network this corresponds to:

- *Observe*: Establish what is where (picture compilation)
- *Orient*: Establish what is happening (situation assessment)
- *Decide*: Establish what should be done about it (threat assessment, planning)
- *Act*: Provide the effect (resource allocation)

In more detail, the OODA loop operates at many levels within the command chain and the required timeline and the available inputs will vary. Moreover, NEC implies achieving effective informed action in environments in which: control is distributed; uncertainty, ambiguity, imprecision and bias are endemic; multiple stakeholders with different aims and objectives are present; and resources are limited and continually vary during the system's operation. To achieve this, it is proposed that such systems be composed of autonomous, reactive and proactive components (henceforth called *agents*) that can sense act and interact in order to achieve individual and collective aims [2].

The system architecture solutions must also exhibit other desirable system attributes if the benefits of NEC are to be lasting. As we desire to connect more and more of the battle space together, this places a *scaleability* requirement upon any architectural approach, one that can withstand, without degradation, the increasing communication and processor loads that will result. In a similar fashion, the elements that we will seek to connect are heterogeneous and our architectural approach must have the *flexibility* to incorporate a wide range of possible legacy systems, often across coalitions, and still deliver the operational performance required. The operational environment is disruptive and we require that the architecture is *robust* in that it will degrade gracefully and predictably, in the face of such forces. In this vein, we also seek *stability*, *security* and *timeliness*. There is, therefore, great interest in how to design systems to deliver NEC with all these desirable properties.

Now, for many of the command and control systems that are currently produced, the emphasis has been on track fusion or picture compilation. In such systems, inputs are well defined, as are the outputs in terms of an accurate and complete tactical picture. Nonetheless mature solutions do not readily exist for situations with high numbers and diversity of targets and a significant level of background noise and clutter (such as in the land and surface domains). Here there is continued reliance on the operator to do part of the task. Within tactical command and control systems, situation assessment has to undertake: the fusion of non-sensory and prior information with the sensor derived inputs, the aggregation of entities within the tactical picture while coping with the inherent uncertainty in the systems and an ever increasing environmental complexity. In addition to all of this, it must provide information in a form that supports the decision-making process appropriate to the level and the tempo of the task.

At other levels, industry is interested in applying information fusion techniques to solve the growing problems in the area of intelligence processing and assessment to support planning and targeting. This is an area where the reliance on the human operator is very significant. Consequently, with the dramatic growth in the volume and range of information sources, the ability of operators to process such information in a timely way is becoming problematic. To overcome this, it may not be desirable to remove the human completely from the process but to 'do more with fewer people' (possibly more reliably) through the appropriate use of automated systems and decision aids.

At the level of an individual platform, the optimal control of sensors with varying levels of autonomous adaptive control is a significant undertaking. For a networked environment (i.e. an environment involving many agents interlinked via some data or communication network) this is made much more complex. However, the advantages to be gained from such networked, or *network enabled,* systems are potentially so significant that this area is progressively becoming a very active research area.

Given this context, the military problems of interest can be characterised by:

- *Wide range of information sources.* Some may be well defined, highly structured and formatted digital data streams, but key information may also be contained in relatively unstructured text reports. There is also a high reliance on imagery and the associated processing. Historical data may also be an important source. Land battle space provides further challenges with heightened difficulties in determining unambiguously what and where the threats are. There are likely to be lots of sources of information and the need to operate at a high tempo. Intelligence, by comparison, can operate at a slower tempo. This brings the need to assess hugely diverse sources of information (both qualitative as well as quantitative) with different time constants and levels of uncertainty.

- *The analysis process is open-ended.* It is not always possible to be precise at the outset about what the analyst is looking for. Equally it is not easy to determine completeness criteria. Therefore, this gives rise to issues such as opportunism, planned perception and planned resolution of ambiguity.

- *Time is a key factor.* This is both in terms of timeliness of results and in terms of the dynamics of evolving situations. Here we are typically seeking both an assessment of current state plus a prediction of what may happen quickly enough such that the assessment is pertinent. This requires us to balance accuracy against speed.

- *The outputs are diverse.* These range from identification of specific threats to trends and relationships.

This is against a back-drop where:

- Many events are occurring simultaneously with differing priorities for attention and need for response.

- The situation is continuously evolving from a strategic viewpoint and hence associated priorities and decisions need to be revisited

- We are operating in a non-cooperative environment that is jamming and disrupting communications between assets as a result of both uncoordinated or co-ordinated attacks

Given this background, technologies are needed that can deliver the following functionality for NEC:

- *Coordinate different assets under incomplete information.* For example, in naval littoral environment scenarios, where we are operating with coalition forces of non-homogeneous capability and where not all platforms may be fully active participants in the network.

- *Distributed decision-making and control between assets. In the same example,* the operation and management of assets local to the platform must adapt in the presence of the shared networked picture.

- *Decentralised information fusion at many levels.* The devolution of the fusion process through out the whole network in a way such that any single fusion component is consistent with an other. This can only be done in a limited fashion at present [4].

- *Flexibility of response.* The ability to function in different environments and operational circumstance. The ability to function under different degrees of control or conversely autonomy is a prerequisite.

The integration of such functionality must support the development of architectural solutions that provide:

- *Scaleable solutions.* That is, the capability to support deployment of large numbers resources at a level that matches the demand without incurring computational or significant communication bandwidth overheads

- *Robust solutions*. The capability to perform effectively despite degradation of the deployed system as a result of environmental damage.

- *Legacy Compatibility.* The capability to accommodate and integrate with the very many legacy systems that has previously been developed in this domain.

- *Decision making traceability.* In the future, there is going to be much more emphasis on traceability of decisions because of increased levels of autonomy. Issues like being able to prove that all the available, relevant information was taken into consideration might be very important. Also the need to make sure that single sourced items of information are not given undue weight, because they get multiply counted without it being obvious.

- *Provide information management and shared awareness across the battle space.* Here a mechanism is needed for controlling the trade-off between having everything right down to the lowest level sensor data available everywhere to much higher level abstractions.

The identification of these challenges has led to the recognition of the need for fundamental research development and integration of technologies. Elements of this challenge will now be tackled in a major new project on Decentralised Data and Information Systems (DDIS) that is funded under the auspices of the BAE Systems / EPSRC strategic partnership agreement and will be coordinated with BAE Systems' other research activities.

2.0 AUTONOMOUS AGENT TECHNOLOGIES

The DDIS approach is built upon the conceptual framework of software agents --- autonomous, flexible problem solvers that interact in flexible ways in order to achieve their individual and collective aims. Specifically, we believe that this autonomy and flexibility is the key to delivering technical solutions for NEC. In more detail, the agents do this by providing mediation and services within OODA loops at all system levels. To be effective in such challenging environments, the agents need to be able to make the best use of available information, be flexible and agile in their decision-making, be aware of other agents, and be adaptive to their changing environment. These agents provide services among themselves and for human operators that:

- *Co-ordinate different assets under incomplete information*. This will contribute, through a robust and consistent approach to uncertainty modelling, towards an enhanced and substantially automated information fusion in Command and Control,

- *Distribute decision-making and control between assets.* In most NEC systems there is no single locus of control and there are multiple stakeholders each with their aims and objectives. Thus, by definition, the decision making process operates under a decentralised control regime and the various assets are under the control of these distinct stakeholders.

- *Decentralised information fusion at many levels.* This can only be done in a somewhat limited fashion at present. Communities of agents, each operating from their own perspective and with their own resources, provide a natural framework to extend this capability to different forms of information.

- *Provide flexibility of response.* For example. Mechanism Design (see below) allows higher-level goals (Command) to be achieved but without centralised control, achieving a level of decentralised autonomy.

In more detail, recent advances in the understanding of the rules of interplay among the agents and the design of these rules have provided an important insight into how we can go about delivering effective NEC. Specifically the interactions of the various autonomous agents within the system, that are necessary in order to achieve their individual and collective aims, can be analysed and designed using techniques from Game Theory (in general) and Mechanism Design (in particular). Game theory is chosen because it has developed powerful tools for analysing decision making in decentralised open systems with multiple autonomous agents. Recently, these tools have been tailored to computational settings to provide a principled foundation for building multiple agent systems. This tailoring gives rise to the field of *computational mechanism design* [3], which applies economic bargaining or bidding principles to computer systems design. Now while mechanism design shows significance promise, to make it applicable in the NEC context, a number of extensions need to be made (see section 3 for more details).

Communication efficacy also raises another important issue. Current decentralised fusion methods require that uncertainty distributions be passed from agent to agent [4]. However for complex distributions, the bandwidth and time constraints will prevent this. Thus a system offering adequate performance within these constraints will require a reduced, but well principled, representation of the uncertainty distribution to be transmitted.

3.0 THE DDIS RESEARCH PROGRAMME

The DDIS project aims to develop techniques, methods and architectures for modelling, designing and building decentralised systems that can bring together information from a variety of heterogeneous sources in order to take informed actions. To do this, the project needs to take a total system view on information and knowledge fusion and to consider the feedback between sensing, decision-making and acting in such systems (as argued in section 1). Moreover, it must be able to achieve these objectives in environments in which: control is distributed; uncertainty, ambiguity, imprecision and bias are endemic; multiple stakeholders with different aims and objectives are present; and resources are limited and continually vary during the system's operation.

To achieve these ambitious aims, we view such systems as being composed of autonomous, reactive and proactive agents that can sense, act and interact in order to achieve individual and collective aims (see section 2). To be effective in such challenging environments, the 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. Thus we need to bring together work in a number of traditionally distinct fields such as information fusion, inference, decision-making and machine learning. Moreover, such agents will invariably need to interact to manage their interdependencies. Such interactions will also need to be highly flexible because of the many environmental uncertainties and changes. Again this requires the synergistic combination of distinct fields including multi-agent systems, game theory, mechanism design and mathematical modelling of collective behaviour.

The intellectual challenges of bringing together these hitherto largely independent research themes (at both the individual and the multiple agent level) are significant. Traditionally, each of these themes has dealt with some particular aspect related to how a system (involving many agents) may learn and reason about its environment allowing it to sense, plan and act in pursuit of the goals and objectives it has been given. However, in this project we seek to bring these components together to produce a coherent whole. Moreover, we seek to do so in a *principled manner* so that we have the best possible idea of how the overall system will behave (when designing it) and maximum confidence that it will achieve its objectives (when it is operational). In undertaking this foundational research, we need to be mindful of the total systems view so that the related components can be combined into a resulting system that is robust, flexible and scalable. This concern means that we will also undertake research into architectures for decentralised data and information systems that exhibit these properties. Such decentralised architectures are crucial because they do not have a single point of failure, there is no central communication bottleneck, and they are the most natural representation for most complex problems. Here the key issues are achieving effective and coherent operation of the many diverse agents in highly challenging domains, while ensuring that the desired non-functional system attributes are attained. Finally, to provide a focus for this integrated view, the ideas and technologies developed within the research programme will be exercised within the exemplar domain of *disaster recovery*. This domain has been chosen since it requires timely decision making and actions in the highly uncertain and dynamic situations highlighted earlier, because it is an important domain in itself, because it is demanding both from a functional and an integrated system point of view and key issues map on to NEC requirements.

The development of decentralised data and information systems that can operate effectively in highly uncertain and dynamic environments is a major research challenge for computer scientists and a key requirement for many industrial and commercial organisations. Moreover, as ever more information sources become available (through the Web, intranets, and the like) the network enabled capability of obtaining and fusing the right information when making decisions and taking actions is becoming increasingly pressing. This problem is exacerbated by the fact that these systems are inherently *open* [5] and need to respond in an *agile* fashion to unpredictable events. Openness, in this context, primarily means that the various agents are owned by a variety of different stakeholders (with their own aims and objectives) and that the set of agents present in the system at any one time varies unpredictably. This, in turn, necessitates a decentralised approach and means that the uncertainty, ambiguity, imprecision and biases that are inherent in the problem are further accentuated. Agility is important because it is often impossible to determine *a priori* exactly what events need to be dealt with, what resources are available, and what actions can be taken.

Clearly, such data and information systems are significantly different in nature from the traditional information systems that are found in many organisations today and they have much more in common with planned developments in the Web, the Grid, Peer-to-Peer Systems, Autonomic Systems, and Pervasive Computing. For this reason, research in this area needs to cover a broad set of issues and topics (hence the breadth and the scale of the project) and needs to be undertaken at many different levels (hence the joint academic/industrial nature of the partnership).

Firstly, there is the *fundamental research* needed on both the individual agents and collections of them. Such research is needed to develop individual agents that can act and interact to achieve their goals in an autonomous fashion by utilising feedback between sensing, decision making and acting. All of these individual activities need to be undertaken in the context of a multiple agent system (i.e. we are concerned with social agents, not solipsistic ones, and so sensing, information fusion, decision making and acting all need to take this into account). In addition to this social influence on their individual actions, the agents need to be able to explicitly initiate and respond to social interactions in order to further their individual and collective aims. Here the challenge is to do so in an agile manner that is robust to the inherent uncertainty, dynamism and resource variability.

Secondly, there is *systems research* in that the methods and techniques emerging from this fundamental research need to be incorporated into DDIS architectures that exhibit the desired functional and non-functional characteristics in the types of environment we consider.

Finally, there is more *applied research* involved in building the software demonstrations. These will help illustrate the potential, the cohesiveness, and the effectiveness of the various techniques and the architectures that have been developed (both in a practical and a quantitative way). They will also provide feedback into the fundamental research by highlighting any shortcomings (in terms of performance or unrealistic assumptions) of the methods that have been developed.

Naturally, there are many interdependencies and interrelationships between these research themes (as shown in Figure 1 below). In particular, there is a strong relationship between the research on individual and multiple agents (since the former are the basic building blocks of the latter and the latter has a profound impact on the behaviour of the former) and the work on demonstrators will, in turn, significantly affect the fundamental research in terms of method appropriateness and performance. Nevertheless, for the purposes of structuring the research programme, we discuss them separately.

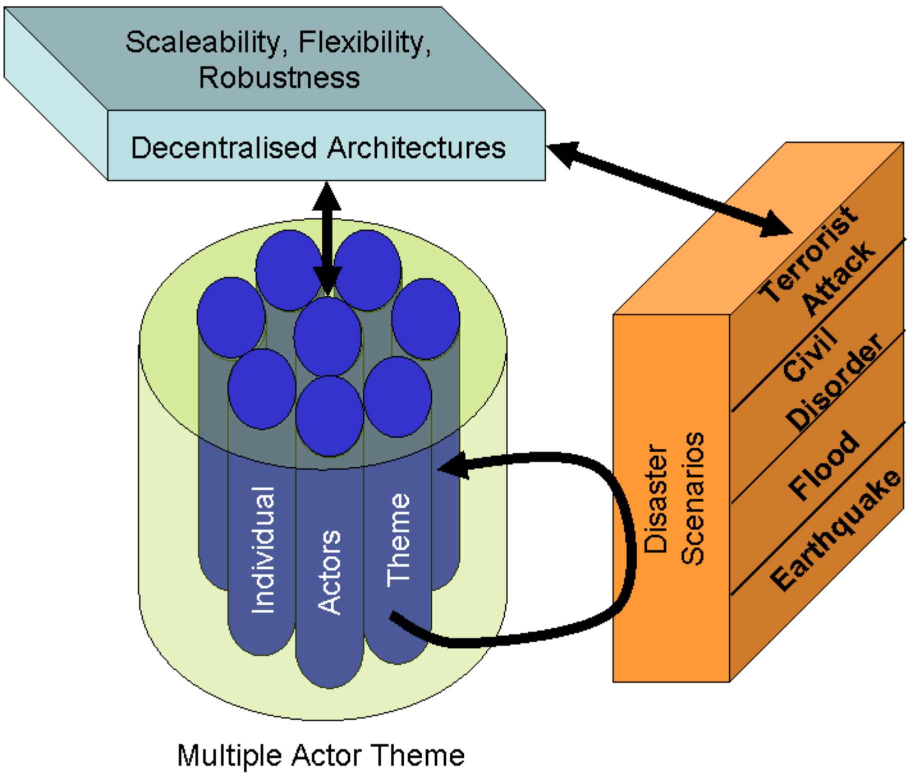

Figure 1: System view of the research programme

3.1 Individual Agents

This research theme is concerned with techniques and methods for designing and developing the individual agents that form the basic building blocks of the DDIS. This is a significant challenge for two main reasons. First, we need to take a holistic view of the individual agent. Thus, each individual 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.

3.2 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 sorts 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.

3.3 Decentralised System Architectures

This research theme is concerned with the study and development of decentralised system architectures that can support the individual and multiple agents in their sensing, decision making and acting in the challenging environments we have previously characterised. The defining characteristic of such systems is that they do not rely on a centralised coordinator or controller. This is motivated both by the inherent structure of the domain/application and a number of perceived system or operational benefits (including fault-tolerance, modularity, scalability, and system flexibility – as detailed in section 1). Now in contrast to their centralised counterparts, decentralised data and information systems offer many advantages. In a centralised system, data is communicated to a designated agent where it is fused and, subsequently, decisions are made. The results of the fusion or decision process are then communicated back to the other agents in the system. However this leaves the system open to many vulnerabilities since the central agent is a single point of failure. Further, such systems place large demands on communications and this limits the size of the system that can be developed.

Given this context, the key research activities involved in this area are:

- to determine the range of issues and variables that will govern the possible architectures and determine how these options can be compared and contrasted;
- to evaluate these options to determine their relative merits in varying circumstances.

3.4 Applications: Disaster Recovery

To ensure the specific methods and techniques developed in the research fit together to give a coherent whole, the project will develop a number of *software demonstrations*. These will be in the broad area of disaster recovery (for the aforementioned reasons) and will focus on specific vignettes. These vignettes will be chosen to fully exercise both the functional and the non-functional requirements of the system.

4.0 EXPLOITATION

The research programme aims to progress algorithms and methods in a number of distinct fields and to integrate them within decentralised architectures with desirable global properties. The global properties of interest are driven by real operational challenges and the success of the algorithms and their integration will be exercised within the applications research theme across a number of disaster recovery vignettes. The vignettes have been chosen to fully exercise the functional and non-functional requirements on systems of interest, are readily understood, are not specific to the interests of specific industrial players and provide a *lingua franca* for would be exploiters to whom they can be made public. This will make it easier for them to identify exploitation opportunities (through the steering committee and events such as the application symposium). The public domain vignettes can be mapped to their own market challenges.

Industry is directly involved in the research, in particular in the Architecture and Applications research themes, and has been actively involved in the shaping of the programme as a whole. There has been a high level of industrial involvement in establishing the aims and objectives of the research programme, and in its shaping to facilitate delivery of viable 'mature' technologies. The output of this research will allow us to understand better the capabilities of different approaches and so more readily meet the future requirements in the NEC environment

Parallel industry activities are recognised as necessary to ensure the transfer to industry and successful exploitation of the research. Some of these activities have already been identified and planned. The principal activity is the exercise of military vignettes to facilitate technology maturity assessment within industry and government, as well as providing feedback to the researchers. Others are yet to be planned in detail, as they will take place in later years and will be dependent on both the research outputs themselves as well as the particular interests of the industrial partners. It is also likely that opportunities will exist to second a number of industrial partner staff to work alongside the researchers to further technology take-up.

5.0 CONCLUSION

The changes in military capability that are required to face the challenges of complex, non-traditional operations drive the need to provide greater flexibility and to respond more quickly to new information. This applies at the strategic, operational and tactical levels. The MoD description of NEC is an important statement of the challenge and begins to shape the requirements for future systems. This situation requires us to invent and develop new and innovative technical solutions. Some of these solutions have been described here.

The analysis presented in this paper of the MoD's NEC requirements has allowed us to identify some key technologies for further research. An important component of this is the derived requirements for decentralised control, autonomous problem solvers and flexible interactions. These requirements can be met by adopting an agent-based approach and there are some tools available for doing this. However, further work is needed to reach the full potential of this approach. Hence the investment in the DDIS programme and the collaboration in shaping it to ensure that it will deliver the solutions required. To do this it is essential to have derived the programme from an understanding of the whole system solutions that industry is best placed to supply together with the innovation for the research base.

This programme is now in place and we look forward to presenting some of the outcomes in future meetings.

References

[1] MoD (2003) “Delivering Security in a Changing World”, Defence White Paper, http://www.mod.uk/publications/whitepaper2003/.

[2] N. R. Jennings (2001) “An agent-based approach for building complex software systems” Comms. of the ACM 44 (4) 35-41.

[3] R. K. Dash, D. C. Parkes and N. R. Jennings (2003) “Computational Mechanism Design: A Call to Arms” IEEE Intelligent Systems 18 (6) 40-47.

[4] P. Vangasse and D. Nicholson (2000) “Architectures for efficient data fusion” *Proc. NATO Advanced Research Workshop on Multisensor Data Fusion*, Pitlochry, Perthshire, Scotland.

[5] C. Hewitt (1991) “Open Information Systems Semantics for distributed artificial intelligence” Artificial Intelligence 47 (1-3) 79-106