Distributed Multi-layered Network Management for NEC Using Multi-Agent Systems

Richard Vaughan, James Wise, Paul Huey, Michael Alcock, Jonathan Vaughan, Steven Shingler, and Graham Atkins

General Dynamics United Kingdom Ltd.
Bryn Brithdir, Oakdale Business Park, Blackwood, Wales, NP12 4AA
richard.vaughan@generaldynamics.uk.com

Abstract. Within the military environment it is important that effective communication lines are maintained so that critical messages are able to reach their destinations and remain secure. Currently the bulk of the communications infrastructure is preplanned prior to mission start, and requires manual intervention when reality fails to match the plan. The communication systems need to be highly flexible, and adaptable in the face of (unforeseen) hostile and adverse conditions.

We believe that this is where a combination of a distributed agent-based system and a reconfigurable peer-to-peer overlay network can be used help to provide a communication system that is robust and highly adaptable in the face of ever changing and adverse conditions with a minimal amount of planning and enable us to reduce the burden placed on dedicated staff in the field.

1 Introduction

During a military mission, the communications network must be structured to support the commander's intent for the mission in all circumstances, which may change several times as the situation unfolds. The combination of constraints (including very low bandwidth channels, high mobility, zero (or minimal) infrastructure and a very adverse environment) and success criteria [1] make this a particularly challenging application domain.

The management of such a network within a dynamic and hostile environment presents a number of problems to which there is no static or single point solution. No control system can accurately and instantly discover the global state of a network, deduce the ideal global state, and enact a change to the ideal state.

We propose to augment the existing planning and management approach by implementing a distributed multi-agent system (MAS) capable of making delegated decisions about the communications infrastructure on behalf of users[2] which will run on the devices in the network.

Section 2 of this paper describes the current approach to management of such a network, section 3 describes our approach to the problem, section 4 describes

[1] i.e. Completion of the mission, irrespective of network performance.
[2] or peer services.

S. Joseph et al. (Eds.): AP2PC 2006, LNAI 4461, pp. 159–166, 2008.

the experimentation environment we have constructed with which to explore the problem space, section 5 defines the military case study within whose boundaries we conduct our experiments, section 6 presents our results and in section 7 we present our conclusions.

2 Management of Existing UK Military Radio Networks

As we discuss in [1], the management of the existing UK military radio networks is a manually intensive, off-line process requiring expert input from signalling staff. The main issues with this management approach are that:

- **Configuration is created off-line prior to mission start.** Flaws in the configuration can lead to invalid communication plans when applied to the devices on the network. These flaws are corrected off-line and the updated plans manually reapplied to the devices.
- **The management applications are separated from the targets of control.** There is a risk that a flaw or network problem that prevents automatic update can can lead to isolated 'islands' which are outside automatic control.
- **Unforeseen changes.** As the situation unfolds on the battlefield an entity might find itself in a position that was not planned to be in. This might mean that it is unable to communicate with other nearby entities as they do not have the correct frequency allocation or cryptographic keys to enable communication. The solution for this requires an element of off-line re-planning and manual re-application of the plans to the devices in the network. However, it may not be possible to get the updated configuration out to the entities in time due to the fluid nature of the battlefield.

3 Our Approach

We propose to augment this planning and management approach by creating a distributed multi-agent system capable of making delegated decisions in a dynamic environment running on the devices in the network rather than simple augmentation of existing logic with agents. The system is capable of determining the appropriate configurations for itself as the system is running, and has the ability to enact the configurations accordingly based on high level plans fed into the system at run-time.

By adding intelligent agents into the system we are able reduce the detail required for the planning stage since the agents will be responsible for the management and implementation of the plan. For example, this means that instead of having to specify which individuals can communicate, the plan can be defined at a higher level so as to identify only the groups or roles that must communicate. Users will be assigned to one or more groups and roles at mission start and identified to the system through their data terminals. The agents are able to fuse this information together at run-time to produce the highly detailed plan and configurations necessary to manage the communications network. This approach offers a great deal of flexibility since it means a user can be removed from one group or role and assigned another while the system is running. The agent system will detect this and be able to reconfigure the communications system accordingly as the (dynamic) policies of the system dictate.

Agents[12], or agent-based systems, are not a panacea but we believe they can provide a more appropriate mapping between the control system and the system being controlled, and allow us to automate many of the resulting control mechanisms. They can also accept delegation of those decision-making and collaborative functions that are more suited to machines than people, easing the burden on operators and providing a management system that can operate in an autonomic[7] manner given appropriate high-level direction[1].

The agent-based system is shielded from the complexities of the underlying system(s) through a layered service architecture capable of presenting information fused from disperate sources and is able to enact the resulting decisions (described further in [1]). The following layers have been identified (see Figure 1):

- An external-facing layer, that will enable external applications to interact with the system as a whole.
- The agent layer itself.
- A P2P layer, which allows the agents to form an overlay network atop the underlying networking and facilitates inter-agent messaging
- A network access layer that allows the agents to influence and retrieve information from the networking sub-system or network management system(s).
- An OS abstraction layer that allows the agent system to interact with the underlying operating system.

We acknowledge that even this MAS will not be able to see the entire network at any moment in time, but, by having the agents know about changes and influences within their local zone[3], we 'believe the agents will be able to make informed decisions more rapidly and efficiently than an equivalent centralised system that can only see traffic passing through, or logged to, its local node.

The agents are governed in their decisions and actions by policies. These policies are updated dynamically as the system runs, as well as being defined up-front during the planning stage(s). Policies will be set by the commanders to inform the agents how their section of the network should be managed to best meet their intent. These policies may be further refined by the agents to direct network or system components under their control or influence as needed.

3.1 System Policies

Policy is defined in the dictionary as: *"A plan or course of action, as of a government, political party, or business, intended to influence and determine decisions, actions, and other matters"*[18]. We have refined this definition to *"an expression of rules intended to govern the behaviour of the system"*.

[3] For example, peers within 2 network hops.

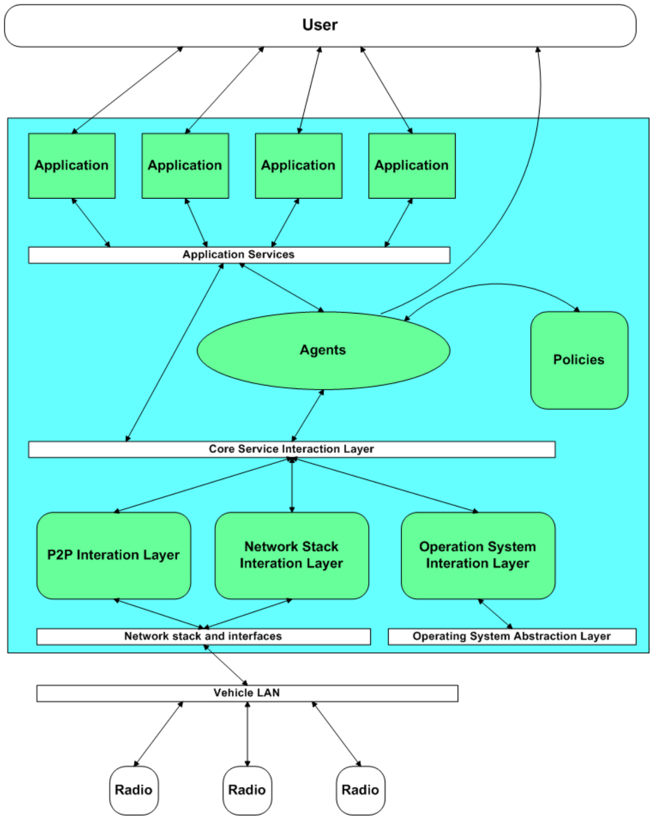

Fig. 1. Application Layers

Having researched recent papers on policy (e.g. [9], [13]), and having looked at some of the available toolkits, we have found no single policy model that embraces all of our requirements and no currently available implementation deals with the adaptive and dynamic nature of our environment appropriately and effectively. This means that we need to extend existing work to create a policy engine and associated verification mechanisms that can operate in our mobile, decentralised environment. Therefore it would appear that the approaches taken by either (or both) Imperial College on the Ponder[10] toolkit, and the work currently underway at DeMontford University on SANTA[11] could be a good base, since, both of these allow for an expression of both management and security policies.

4 Experiment Case Study

Our research has been conducted within the context of the military case study shown in Figure 3: An agile military mission group is moving through an environment where other networks (military or otherwise) are present. The terrain

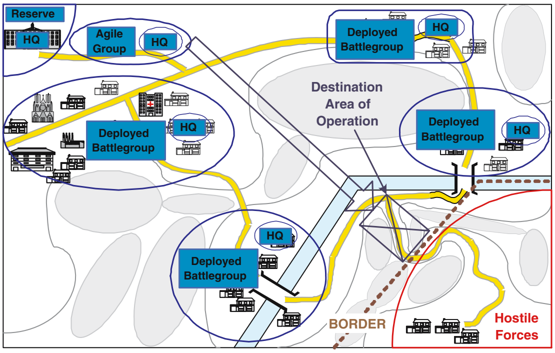

Fig. 2. Basic Mission Scenario

through which the group is travelling contains various features that can block (or disrupt) radio transmissions. The radio network must be managed such that:

1. It is always possible to communicate back to HQ (directly or indirectly).
2. The available radio network is used in the most effective manner.
3. Minor problems with the radio network are detected, analysed, and, fixed.
4. If a fix for the problem is beyond the system's capability, a human is informed to guide corrective action.
5. Most importantly, the current commander's intent for the mission must be supported effectively at all times.

5 Results

Our initial research has led us to implement an initial release of the framework for conducting these experiments and prototypes of the agents and P2P services required by our architecture[1].

We have agents that are capable of detecting simple kinds of interference on the radio network, and of making a decision as to whether the system should change radio channel given the changing nature of the user requirements and environmental conditions. This channel change is propagated across all vehicles considered to be on the local vehicle network.

The agents in the agile group (through the P2P layer) are able to maintain communications with peer agents within the headquarters network. The P2P layer is responsible for the actual maintenance of the network routes and data pipes, but the agents monitor and direct the P2P layer accordingly.

Figure 5 shows how this basic control flow occurs in the system.

1. An application requests access to some resource on another network or vehicle.
2. The local agent talks to the remote agent (via a local P2P service) and requests the service.

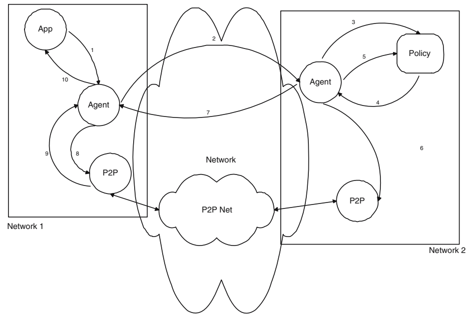

Fig. 3. Basic control flow in the network

3. The remote agent checks the request against his current policies to see if it is allowed.
4. The policies indicate that it is allowed.
5. The remote agent creates a new policy for his P2P service to allow the P2P layer to set the service up.
6. The remote agent informs his local P2P layer about the new policy.
7. The remote agent communicates the access to the local agent.
8. The local agent instructs the local P2P service to setup the user service.
9. The local P2P layer notifies the agent when the user service is ready.
10. The local agent notifies the application that the user service is ready and he can start using it.

The design of our test environment constrains our inter-agent chatter to only that which is necessary, and forces them to use a specific managed pipe (across the P2P overlay network) when it wants to communicate to an agent outside of it's local node rather than allowing the agent-framework to dictate what messages are sent, when, and how. This is important in military networks as the mobile sections are unlikely to have bandwidth to waste, and we do not want an agent to update all of it's peers to the detriment of the networks intended use.

6 Conclusions

We believe that we have defined an authentic case study in which agents can provide a benefit as part of a system level solution. This is important for the agent community, as, without authentic, customer focused, case studies based on domain specific problems, we cannot prove whether solutions incorporating agents or agent-based systems will offer an advantage over those that do not[19].

Along the way we have discovered just how complex distributed multi-agent systems are to implement. Since we have chosen to implement a complete end-to-end system (rather than a section or component of a system) we have had to deal with the complexities of the many interactions between the components that comprise the system and which all need to be defined, implemented, tested, and debugged.

Whilst, at this time we cannot say with any certainty that we have identified a *'killer app'* for agents, we believe that military radio network control systems provide challenges that agent systems are well placed to solve. The flexibility for which deployment of a MAS provides persuasive solutions that are likely to contribute effectively to system level objectives without requiring a mass increase in the workload or expertise of the ordinary users of the communications systems.

The flexibility and complexity of the communications system will increase as new radio technologies promote an expectation of ubiquitous, always-on connectivity and provide a greater variety of connection types and capabilities at each communication node. This particularly applies with reconfigurable software radios ([22], [23]) and the consequent opportunity for selective use of adaptive protocols and policies. At this stage, direct control of radio parameters becomes infeasible, in the same way that pilots of fast jets no longer stabilise the airframe manually when flying.

Our results have shown that agents are well suited to accepting delegation of routine and remedial management tasks currently carried out by signalling staff and are able cooperate to ensure system level changes are monitored and controlled effectively and consistently in a decentralised environment.

Our results also show that the overhead incurred from the distributed solution does not overwhelm or cause the network to be ineffective in normal use. For now this is a binary test of message arrival. We are designing more experiments that will produce more accurate metrics regarding the overheads concerned in benign and emergency use cases. We will present these findings in a later paper.

We are confident that the proposed agent-based approach will enable us to build systems which are more adaptable in the face of change, and that this kind of system provides a sustainable upgrade route, since the agents, and the service layers can be replaced (or augmented) without breaking the overall system [4].

Acknowledgements

This work is supported by the UK MOD Data and Information Fusion Defense Technology Centre.

References

1. Distributed Decision-making and Control for Agile Military Radio Networks In: DAMAS 2005 (2005)
2. Project: Bowman, `http://www.mod.uk/dpa/projects/bowman.htm`
3. Single Channel Ground and Airborne Radio System (SINCGARS), `http://www.fas.org/man/dod-101/sys/land/sincgars.htm`
4. Clansman, `http://www.army.mod.uk/equipment/cs/cs_cln.htm`

[4] As long as appropriate design guidelines are implemented and adhered to!

5. The FreeBSD Project, http://www.freebsd.org
6. IMUNES, An Integrated Multiprotocol Network Emulator/Simulator, http://www.tel.fer.he/imunes/
7. Autonomic Computing, http://www.research.ibm.com/autonomic/
8. Network Coding, http://www.networkcoding.info
9. An Artificial Intelligence Perspective on Autonomic Computing Policies, http://www.research.ibm.com/people/w/wwalsh1/Papers/policy04-acp.pdf
10. Ponder: A Policy Language for Distributed Systems Management, http://www-dse.doc.ic.ac.uk/Research/policies/ponder.shtml
11. Janicke, H., Siewe, F., Jones, K., Cau, A., Zedan, H.: Analysis and Run-time Verification of Dynamic Security Policies. In: Thompson, S.G., Ghanea-Hercock, R. (eds.) DAMAS 2005. LNCS (LNAI), vol. 3890, pp. 92–103. Springer, Heidelberg (2006)
12. Software Agents: An Overview, http://www.sce.carleton.ca/netmanage/docs/AgentsOverview/ao.html
13. Integrating goal specification in policy-based management. In: 2nd International Workshop on Policies for Distributed Systems and Networks (2001)
14. Network Enabled Capability (NEC) - Joint Services Publication (JSP 777), http://www.mod.uk/issues/nec
15. Schumacher, M.: Objective Coordination in Multi-Agent System Engineering. LNCS (LNAI), vol. 2039. Springer, Heidelberg (2001)
16. Vaughan, R.: Agent Based Routing (not yet published)
17. Multilateral Interoperability Programme (MIP), http://www.mip-site.org
18. Definition of policy, http://dictionary.reference.com/search?q=policy
19. Wagner, Gasser, Luck: Impact for Agents. In: AAMAS (2005)
20. Probability Theory: The Logic Of Science E. T. Jaynes (ISBN 0521592712)
21. Reinforcement Learning Repository, http://www-anw.cs.umass.edu/rlr/
22. Joint Tactical Radio System, http://jtrs.army.mil
23. Next Generation Communications (XG), http://www.darpa.mil/ato/programs/XG/