Agents in Logistics Planning – Experiences with the Coalition Agents Experiment Project

Don Perugini [1, 2], Steven Wark [1], Andrew Zschorn [1], Dale Lambert [1], Leon Sterling [2], Adrian Pearce [2]

[1] Defence Science and Technology Organisation (DSTO), PO Box 1500 Edinburgh, South Australia 5108, Australia
{firstname.lastname}@dsto.defence.gov.au

[2] Department of Computer Science and Software Engineering, The University of Melbourne, Victoria 3010
{leon / pearce}@cs.mu.oz.au

Abstract

Military logistics planning is a complex process, involving many calculations, satisfaction of constraints, and cooperation amongst many organisational entities that provide services in order to achieve military logistics goals. Multi-Agent Logistics Tool (MALT) is a project aimed at supporting military logistics planning. MALT is being developed using agent technology, where agents represent the organisations within the logistics domain, and model their logistics functions, processes, expertise, and interactions with other organisations. Agents are a suitable technology for modelling organisations within MALT, due to the similarity in characteristics between organisations and agents. A component of MALT was implemented within DARPA's Coalition Agent Experiment (CoAX) project. We discuss the CoAX implementation of MALT, and lessons learnt. We discovered that implementing a centralised agent planning approach within MALT, and hence the decentralised military (operational) logistics planning domain, may not always be appropriate, and that a decentralised agent planning approach may be more suitable. Some of our observations regarding the future of agents for military logistics planning are discussed.

Keywords: Military Logistics Planning, Decentralised Agent Planning, BDI, CoAX

1. Military Logistics Planning

Military (operational) logistics planning primarily involves supplying and transporting resources and military assets. Logistics planners must form a plan to achieve specific logistics goals, such as deploying force elements[1] and sustaining them throughout an operation (by providing food, water, fuel, ammunition, medical support, storage, etc.), and planning a medical evacuation (*medevac*) when casualties occur. These logistics goals are achieved by obtaining services from various organisational entities, for example, obtaining fuel from a fuel supplier, and having a freight company provide the transportation of the fuel from the fuel supplier to its required destination.

The military logistics planning domain is decentralised because it typically requires the involvement of separate independent organisations – as a combination of Australian Defence Force (ADF), civilian and coalition organisations could be involved. The organisations, which primarily include supply, transport and force element organisations, are geographically distributed, and must cooperate in order to achieve the logistics goals. Each organisation have their own logistics "business" processes in order to perform their particular logistics functions (services), required to achieve logistics goals. The logistics domain is decentralised, not because it is geographically distributed, but because it's organisations exhibit a strong notion of autonomy with characteristics such as: being self interested; making their own decisions (i.e. not controlled by others); and being reluctant to release information (e.g. because it may be proprietary or classified). The logistics domain is also dynamic, where logistics goals, organisations' capabilities (the type and availability of services they can provide) and beliefs are continually changing throughout the planning process; as well as open, where organisations may enter or leave the system at any time.

[1] A Force Element (FE) is a military unit and its associated resources/assets and equipment, including personnel, vehicles, aircraft, and weapons.

To add to the complexity, logistics planning requires many interactions between organisations, many calculations and satisfaction of many constraints (e.g. ensure that casualties are delivered to the appropriate medical facilities in time). As a result of logistics planning complexities, there is typically a trade off between the time to form the logistics plans and the quality of the logistics plans formed. Multi-Agent Logistics Tool (MALT, formerly LPS) [1] is a project aimed towards developing a military logistics planning support system, which will automate aspects of logistics plan formation, analysis and information gathering.

MALT comprises agents that represent the organisations, modelling their logistics business processes, expertise, and interactions with other agents in the logistics domain, in order to achieve the organisations' logistics functions. Agents in MALT cooperate with each other in order to form a distributed logistics plan (services from various organisations) to meet their logistics goals. The input for MALT is a logistics goal, and the output is a logistics plan, which can be executed in order to achieve the logistics goal.

In this paper, we will discuss the reasons for using agents in the logistics domain, our experiences with implementing components of MALT within DARPA's Coalition Agent eXperiment (CoAX) [2] project, lessons learnt in MALT's implementation, and our observations regarding the future of agent technology for logistics.

2. Why use agents?

Agents have strong autonomous characteristics that distinguish them from other software paradigms:

"Objects do it for free; agents do it because they want to." [3].

In addition to autonomy, agents are proactive (goal directed and thus intentional), reactive, and exhibit complex social behaviour (rather than software entities that can execute each other's functions/methods freely) [3].

The logistics domain is distributed and involves decentralized (autonomous) organisations. These organisations are also: intentional entities, with goals (functions and roles), beliefs, and use processes and expertise in order to achieve their goals; are reactive, and thus respond to changes that occur in their environment; and are social, so they interact with other organisations to achieve their goals, where the social interaction is typically complex, such as negotiation, rather than just action requests. The similarity in characteristics between agents and organisations make agents an appropriate choice for modelling organisations within MALT.

3. MALT within CoAX

DARPA's Coalition Agent eXperiment (CoAX) [2] project was aimed at demonstrating the utility of agents for coalition planning [2, 4]. Some 20 organisations from the USA, UK and Australia were involved. MALT was implemented within CoAX (Oct 2002) [4], with the aims of demonstrating components of MALT, and thus the use of agents for military logistics planning in the Australian-Coalition contexts, interoperability with foreign agents (i.e. agents not developed by us), and the dynamic capability of MALT.

Our part of the CoAX project involved a vignette where an Australian ship was struck by a torpedo, resulting in damage to the ship and casualties. A medevac is required to transport the casualties from the ship to a coalition medical facility using available helicopters. Agents represented the ship's medevac function (*medevac agent*) and onboard resources, and a single proxy agent represented the coalition helicopter and medical facility resources, providing information regarding the availability of medical facilities to treat casualties, and availability of helicopters to transport the casualties. The medevac agent modelled the logistics process of planning a medical evacuation from the ship. The agents cooperated in order to form a medevac logistics plan.

The medevac agent was developed using the ATTITUDE multi-agent architecture [5, 6], which is based on the belief-desire-intention (BDI) agent model [7]. The logistics process for a medevac was represented as a plan (routine), and executed when casualties were detected on the Australian ship, triggering the "medevac" task. Cooperation between the medevac agent and the other agents were facilitated by DARPA's CoABS grid [8] and IHMC's KAoS agent management system [9].

The medevac agent, when triggered, requests the availability of medical facilities to treat its casualties, and helicopter resources to plan the transportation of casualties from the Australian ship to the nearest suitable medical facility. The proxy agent responds with available medical facilities and helicopters, providing distances to facilities, start location, earliest start time, and types of helicopters available. The medevac agent uses prior knowledge of the carrying capacity and speed of the types of helicopters.

A simple algorithm is used by the medevac agent to form a transportation plan. The helicopter that can transport the injured to the medical facility at the earliest time is selected to perform the transportation task. Highest priority casualties are transported first. If the selected helicopter cannot transport all the injured, the process is continued with the remaining injured.

The plan formed is sent to the foreign Multi-Level Coordination Agent for processing, to deconflict and optimise (merge) the medevac plan with existing flight plans developed by foreign (coalition) agents. The plan is then distributed to the appropriate coalition helicopters and medical facility for execution. The medevac agent reacts to any changes, such as helicopter availability, the number and type of casualties, or the availability of the helicopter landing pads and replans if necessary.

The CoAX demonstration was held in October 2002 at the US Navy Warfare College, Newport RI. It successfully presented components of MALT in operation, and thus demonstrated the use of agent technology for military logistics planning. UltraLog (or ALP) [10] has also demonstrated that agents can be used for the U.S. logistics planning domain. We are focusing on the Australian-Coalition contexts. The logistics process of a (ship) medevac was effectively modelled in the medevac agent using ATTITUDE. The medevac agent successfully cooperated with foreign agents, sending the final medevac plan to the foreign Multi-Level Coordination Agent for processing and distribution to coalition helicopters and medical facility. The medevac agent was able to dynamically replan when the situation changed.

4. Lessons Learnt

This architecture for logistics planning was used in the CoAX context to facilitate development and testing of the Australian agents independently from the rest of the coalition organisation. Although this transportation planning approach worked well for the CoAX demonstration, its essentially *centralised* nature imposed several limitations. By centralised, we mean that all information about the helicopters agent's capabilities are sent and processed by the centralised medevac agent, which forms a transportation plan by deciding (alone) how the coalition helicopters' capabilities will be used to achieve its transportation task. This may be appropriate if the medevac agent and helicopters are from the same organisation (e.g. ADF), which typically act towards a common goal and may have a limited right to issue orders among themselves. However, if other organisations are involved, e.g. coalition organisations, such as the coalition helicopters, this may not always be appropriate. Coalition helicopters, which are from one organisation, may not want to be told what to do by the ADF medevac agent that is from another organisation. The coalition helicopters may be self-interested and decisions made by the medevac agent may not be in the interests of the coalition helicopters. Also, information by one organisation (e.g. coalition helicopters) may not be able to be released for processing by a centralised agent that is from another organisation (e.g. ADF medevac agent), because the information may be proprietary or classified. For example, in the CoAX experiment, coalition helicopters were not immediately available as they had already been scheduled for other tasks – information not released to the medevac agent. As a result, the Multi-level Coordination Agent described above was needed to resolve any conflicts and maximise synergies between the medevac plan and the existing coalition flight plans. In a decentralised system this would be embedded within the negotiation process.

Limitations of a centralised approach are not only confined to coalition operations. Most ADF logistics operations are likely to involve civilian (other) organisations. Therefore, in general, a centralised agent planning approach may not always be appropriate within MALT, and hence the decentralised military (operational) logistics planning domain.

The centralised agent planning approach also has a technical limitation. The medevac agent used a simplified protocol that assumed that the helicopters are available anytime after its specified start time, and can carry its full capacity of casualties throughout the time it is available. If the medevac agent was to accommodate situations such as helicopters not being available at various times after their specified start time, or can arrive at the ship during some other trip (piggyback), and hence can only carry a portion of its full capacity, then the coalition agent will have to communicate much more information, possibly the complete flight plans of all helicopters in the theatre of operations, to the centralised medevac agent. The quantity of information that would need to be communicated to the medevac agent in such a case would be extensive. If a decentralised approach was implemented, agents could process the information themselves and send just the results.

A new protocol is being developed, called the Provisional Agreement Protocol (PAP), which allows decentralised agent planning [11, 12]. An agent's logistics goal is sent to service providing agents that may be used to achieve the goal. The service providing agents only release capabilities (services) that they are willing to perform in order to fully, or partially, achieve the logistics goal. The most suited capability is selected. If the selected capability did not achieve the complete logistics goal, the portion of the logistics goal that the capability did not achieve becomes the new logistics goal to be achieved. The process repeats until the logistics goal is completely achieved, resulting in a distributed plan that describes the capabilities to be performed by the various service providing agents in order to achieve the logistics goal. The protocol allows backtracking, hence if a selected capability was later found to be inappropriate to achieve the logistics goal, then the agent performing the planning process may de-commit from the selected capability, and select another capability to replace it. The protocol provides policies regarding the commitment and persistence of agents' capabilities and goals, and speech acts to facilitate this interaction (planning process).

5. Future of Agents for Logistics

Based on our experiences with using agents for logistics planning, we found two main issues[2] that need to be overcome in order for agents to be effectively used for military logistics planning – technology and social (human) acceptability.

5.1 Technology

- ***Logistics business process modelling*** – How can the logistics processes be effectively modelled using agents? A framework is required that allows a developer to extract logistics processes from an organisation that are to be automated, and embed these processes into an agent(s). The framework will most likely be dependent on the agent architecture, e.g. BDI agent architecture. If the organisations logistics processes are dynamic, how can they be maintained in the agent system? There are also technical issues that need to be resolved regarding logistics processes, such as how to cope with late or inconsistent information, and how can learning be used?
- ***Protocols*** – Need to devise protocols that can facilitate decentralised agent planning, that will consider the type of interactions and negotiations that take place between organisations within the logistics (business/e-commerce) domain. This work is currently under development.
- ***Ontologies*** – What knowledge is required for the logistics domain, and how does one effectively represent it and reason about it? Ontologies will allow agent interactions by providing agents with common semantics, enabling agents to understand communicated information from other agents, such as logistics goals, services (capabilities), and descriptions of resources (physical objects). The type of reasoning that can be achieved, based on requirements, needs to be investigated, since there is a trade off between expressiveness and computational complexity. CoAX demonstrated the use of DAML, *"which is a language that provides a rich set of constructs with which to create ontologies and to markup information so that it is machine readable and understandable"* (*www.daml.org*).
- ***Automated information gathering*** – MALT intends to incorporate agent "experts" who can provide advice and knowledge to other agents based on their domain of expertise. Expert system technology typically focuses on the single agent. Complex social issues need to be investigated, such as providing a mechanism for agents to find the appropriate expertise, particularly when the expertise may be the conjunction of knowledge provided by more than one expert agent. Agents may search for information themselves, where the information is stored in various information sources. This will be facilitated using semantic web technologies, such as DAML, where information is tagged with knowledge, based on some ontology, making the information "machine readable", allowing agents to search for information on their own (in an automated way). Elements of automated information gathering were demonstrated in CoAX, including the briefing agent by DSTO [4], Mobile Agent for Medical Monitoring by Dartmouth College (*agents.cs.Dartmouth.edu*), Verona by

---

[2] This list is by no means complete.

GITI and ISX (*coabs.globalinfotek.com*), Decision Desktop by QinetiQ (*www.qinetiq.com*), and Ariadne by University of Southern California and ISI (*www.isi.edu/ariadne*).

- ***Security*** – Security has traditionally been a major concern with information technology systems. Issues such as integrity, authentication, and secure communication, need to be considered. In CoAX, policy control of agent registration, behaviour and communication was demonstrated using the KAoS system developed by IHMC [9].

5.2 Social Acceptability

- ***Trusting agents to do business for you*** – Will humans or organisations trust agents to do business for them without supervision, i.e. have agents sell and buy services and products? There may be a concern that software may not act rationally, and result in the organisation losing money, or possibly worse consequences. Can we ensure that agents will always do the correct thing, or can we demonstrate that agents will do the correct thing most of the time (everyone makes mistakes, even humans), in order to increase the human level of trust in agents? Can we implement mechanisms that can prevent seriously damaging actions by agents? Our initial implementation of MALT intends to form a logistics plan but not act on any of the elements in the plan. The plan is intended as a possible course of action, and if the user finds that it is suitable, the user may act on it. Whether we will get to a stage where the user will allow the agent to act will remain to be seen.
- ***Accountability and the law*** – If an agent does something wrong, who will be accountable? Is it the organisation or the agent developer?
- ***Humans and agents working together*** – Rather than having humans "using" the system, we may want humans as "part of" the agent system. How can this be done effectively, and what effect on the systems overall performance will it have? Suitable *human friendly* protocols and interfaces need to be in place in order for this to be realised.
- ***Agents can do logistics planning more efficiently*** – It needs to be shown that agents can do aspects of logistics planning more efficiently.
- ***Ease of use*** – We need to bridge the gap between the human and the machine, so that humans can use the system with ease. Such a task may not be trivial, particularly since it may be dependent on the individual user. Hence, an interface that can be adjusted to the users preferences may be required.
- ***Adjustable autonomy*** *[13]* – The level of control that the user has over the agents (or system) should be adjustable. Some users may want the agent to make all the decisions (press a button and get all the answers), where others may want to take more control over the agent's actions, specifying exactly how the agent should solve the particular problem.
- ***Adjustable visibility*** – The level of visibility of what the agent does and how it came to a decision should be adjustable. Some users may not care how an agent came to a decision, where others may. Higher visibility is useful particularly if the user wants to gain trust in the agent's decisions.
- ***Social Acceptability versus Optimality*** **–** First, humans may not accept a plan even though it may be optimal. For example, it may be cheaper to schedule flights for a husband and wife on separate flights, but the husband and wife may not accept that plan. Second, agents may undergo complex negotiation processes and protocols, possibly allowing them to obtain better results. Humans may not be willing to do the same, because of complexity (hard to understand) or long iterative processes required (may not want to waste time performing some process over and over again), and thus may be willing to sacrifice optimality for simple and quick negotiation strategies.

6. Conclusion

MALT is a military logistics planning support tool, aimed at automating the complex process of plan formation, information gathering and analysis. It is being developed using agents, where agents represent the organisations within the logistics domain, and model their logistics functions, processes, expertise, and interactions with other organisations. Due to the similarity in characteristics between organisations and (intelligent) agents, agent technology is an appropriate choice for modelling organisations in the logistics domain. MALT was implemented within CoAX, and the aims of MALT were achieved. Lessons learnt from the CoAX implementation was that the centralised planning approach used may not be appropriate for MALT, and hence the decentralised military logistics planning domain, and that a decentralised agent planning approach may be more suited. A protocol, called Provisional Agreement Protocol (PAP) is being developed to facilitate decentralised agent planning. Finally, a list of issues regarding the future of agents for logistics planning was discussed, comprising issues relating to technology and social agreement.

7. References

1. Perugini, D., et al. *Agents for Military Logistic Planning*. in *Workshop "Agent Technology in Logistics" as part of the 15th European Conference on Artificial Intelligence (ECAI-2002)*. 2002. Lyon, France.
2. Allsopp, D.N. et al., *Coalition Agents Experiment: Multiagent Cooperation In International Coalitions.* IEEE Intelligent Systems (see also http://www.aiai.ed.ac.uk/project/coax/), 2002(May/June 2002): p. 26-35.
3. Wooldridge, M., *An Introduction to Multi-Agent Systems*. 2002: John Wiley & Sons.
4. Wark, S., et al. *Dynamic Agent Systems in the CoAX Binni 2002 Experiment*. in *Proceedings of the 6th International Conference on Information Fusion*. 2003. Cairns, Australia.
5. Lambert, D. *Advisers with Attitude for Situation Awareness*. in *Proceedings of the 1999 Workshop on Defense Applications of Signal Processing*. 1999. LaSalle, Illinois.
6. Lambert, D. *Automating Cognitive Routines*. in *Proceedings of the 6th International Conference on Information Fusion*. 2003. Cairns, Australia.
7. Bratman, M.E., D.J. Israel, and M.E. Pollack, *Plans and Resource-Bounded Practical Reasoning*. Computational Intelligence, 1988. **4**(4): p. 349-355.
8. Kettler, B.P. (ed.), *The CoABS Grid: Technical Vision, http://coabs.globalinfotek.com/*.
9. Bradshaw, J.M., et. al., *Making agents acceptable to people*, ed. N.Z.J. Liu. 2003, Amsterdam, Holland: IOS Press.
10. *Ultra*Log Homepage, Defense Advanced Research Projects Agency,*. www.darpa.mil/tto/programs/ultraLog.html.
11. Perugini, D., et al. *Distributed Information Fusion Agents*. in *Proceedings of the 6th International Conference on Information Fusion*. 2003. Cairns, Australia.
12. Perugini, D., et al. *A Distributed Agent Approach to Global Transportation Scheduling*. in *IEEE/WIC International Conference on Intelligent Agent Technology (IAT 2003)*. 2003. Halifax, Canada.
13. Scerri, P., D. Pynadath, and M. Tambe. *Adjustable autonomy in real-world multi-agent environments*. in *International Conference on Autonomous Agents (Agents'01)*. 2001.