Coordination in Distributed Intelligent Systems Applications

Abdeslem Boukhtouta[1,4], Jean Berger[1], Ranjeev Mittu[2], Abdellah Bedrouni[3]

[1] Defence Research & Development Canada – Valcartier, 2459 Pie-XI Blvd. North, Val-Belair, G3J 1X5 (Quebec), Canada
[2] US Naval Research Laboratory, Washington, 4555 Overlook Avenue, SW, Washington, DC 20375, USA
[3] Concordia University, Department of Industrial Engineering, 1455 de Maisonneuve, Montreal, Quebec, H3G 1M8, Canada
[4] Concordia University, Concordia Institute for Information Systems Engineering (CIISE), 1455 de Maisonneuve, Montreal, Quebec, H3G 1M8, Canada

Abstract This paper is devoted to describing the broad range application domains which implement many of the coordination strategies and techniques from the field of multi-agent systems. The domains include defense, transportation, health care, telecommunication and e-business. The purpose of this paper is to describe the diversity of the applications in which multi-agent coordination techniques have been applied to overcome the challenges or obstacles that have existed with regard to performance, interoperability and/or scalability. While the number of application domains is steadily increasing, the intent of this paper is to provide a small sampling of domains which are applying coordination techniques to build intelligent systems. This paper will also describe an emerging and important problem domain which requires the coordination among many entities across the civil-military boundary, and can benefit from multi-agent coordination techniques.

1 Introduction

It is commonly recognized that many different disciplines can contribute to a better understanding of the coordination function as a way to build and provide appropriate tools and adequate approaches for designing complex organizations and systems. Agents incorporated into these self-regulating entities represent "communities of concurrent processes which, in their interactions, strategies, and competition for resources, behave like whole ecologies". We describe in this paper applications requiring coordination or distributed decision-making. A common characteristic of these applications is the highly distributed nature of them. These applications can be also seen as challenging problems for multi-agent coordination (See Panait and Luke [0]).

The research community working in the area of Distributed Artificial Intelligence (DAI) unanimously endorses the idea that coordination – a fundamental *paradigm* – represents a challenging ***research area***. A clear consensus has thus emerged around a forged, well-articulated, and strongly advocated common vision that coordination is a central issue to agent-based systems engineering research (See Bedrouni et al. [2] for more details).

This paper provides an overview of the application domains in which multi-agent systems have been developed. The inherent distributed nature of these application domains reveals that benefits in performance or efficiency can be derived through the coordination between multiple agents. In other words, the collective behaviour of the system is improved through the coordinated interaction of their parts. In this paper, we will cover applications from defence, transportation, health care, telecommunication and e-business.

From the defence sector, we will describe how multi-agent systems have been deployed to support course of action analysis through their ability to monitor the battlefield environment as depicted in command and control systems, and additionally by tapping into simulations as ground truth are able to reason about deviations occurring in the movement of entities represented in the command and control systems. We also describe a large multinational experiment in which hundreds of agents were federated in a simulated coalition scenario. From the transportation industry, multi-agent coordination techniques have been applied to manage and improve the overall efficiency of incoming aircraft. From the

health care industry, distributed multi-agent system coordination concepts have been prototyped to support the monitoring and treatment of diabetic patients. In the area of communication networks, coordination techniques have been applied to diagnose faults in such networks through agents that are able to communicate and exchange local information in order to understand non-local problems. In another setting, coordination techniques have been applied to more efficiently route information in a communication network, utilizing auction mechanisms to purchase or sell commodities such as bandwidth on links and paths, comprised of multiple links from source to destination. Lastly, an e-business application is described in the area of supply chain management, which incorporates Coloured Petri nets as a coordination mechanism between agents.

The paper will also describe an emerging problem domain requiring the coordination of large teams that span the civil-military boundary, specifically during stability, security, transition and reconstruction operations. These operations require the coordination of the military with the civilian sector and non-governmental organizations in response to either natural or man-made disasters. The goal of these operations is to bring stability, rebuild or provide security in order to maintain order during such crisis situations. We will describe how multi-agent coordination techniques can provide tremendous value in these situations.

The remainder of this paper is organized as fellows. Section 2 describes broad range of military applications which implement many of the coordination strategies and techniques from the field of multi-agent systems. Section 3 describes a set of problems drawn from transportation domain and using the distributed artificial intelligence paradigm. Section 4 discusses the healthcare applications. On the other hand, Section 5 describes various applications related to the communication networks. Furthermore, Section 6 is specifically devoted to the coordination in E-business applications. Section 7 discusses the coordination for emergency management and disaster relief. Application of the coordination for tired systems is given in Section 8 and for ambient intelligent in Section 9. Finally, some concluding remarks and recommendations are given in Section 10.

2 Defence

2.1 Control of Swarms of Unmanned Aerial Vehicles (UAVs)

The ability to coordinate military operations with the aid of information technology has become an imperative for command and control systems for the past several years. One of the most important military problems (in vogue) where coordination mechanisms should be used is in the control of swarms of UAVs (unmanned aerial vehicles). The UAVs are considered in this case as highly mobile agents that can be used for performing reconnaissance in a combat environment. In this case the communication between the UAVs is too costly because it would reveal their location to the enemy. The UAVs are sent out to collect information from certain locations and then return. The information sought by a UAV may depend on other information that it has collected by another UAV. There are many other cooperative control problems of the UAVS, we enumerate: cooperative target assignment, coordinated UAVs intercept, path planning, feasible trajectory generation, .etc (See [3]). A recent study of performance prediction of an unmanned airborne vehicle multi-agent system has been developed by Lian and Deshmukh [4]. The Markov Decision Processes (MDP) techniques and the Dynamic Programming approach have been widely used to solve the problems cited above (See Goldman and Zilberstein [5]).

2.2 Coordination for Joint Fires Support (JFS)

Coordination techniques are also useful for the systems that address the Joint Fire Support (JFS) problem. The mandate of Joint fires is to assist air, land, maritime, amphibious, and special operations forces to move, manoeuvre, and control territory, populations, and airspace. The JFS and coalition operations require shared approaches and technologies. The challenge is large, however, with a legion of command-and-control computer systems all developed along slightly different lines of approach. Some

of these systems operate jointly within their spheres of influence. The bigger challenge is tying these and other efforts together in a shared global information network, linking ground, air and sea forces.

2.3 Simulation of C4I interoperability

As the complexity of modern warfare increases, managing and interpreting operational data will continue to be one of the greatest challenges to commanders and their staffs. The wealth of data collected and distributed via Command, Control, Communications, Computers and Intelligence (C4I) systems during battlefield operations is staggering. The ability to effectively identify trends in such data, and make predictions on battlefield outcomes in order to affect planning is essential for mission success. Future commanders will be forced to rely upon new information technologies to support them in making decisions.

Simulations have been used by analysis and planning staffs for years during exercises and operations. Typically, combat simulations are used most heavily during the planning stages of an operation, prior to execution. However, simulations are increasingly being used during operations to perform course of action analysis (COAA) and forecast future conditions on the battlefield. Recent efforts by the Defense Modeling and Simulation Office (DMSO) to improve the interoperability of C4I systems with simulations have provided a powerful means for rapid initialization of simulations and analysis during exercises, and have made simulations more responsive and useable during the execution portion of an exercise. Real-time interfaces between C4I systems such as the Global Command and Control System (GCCS), and the Integrated Theater Engagement Model (ITEM) simulation have provided command staffs with the capability to perform faster, more complete, COAA.

As can be seen in Figure 1, intelligent agents, coupled to C4I systems and simulations, offer another technology to help commanders manage information on the battlefield. The Defense Advanced Research Projects Agency (DARPA) has sponsored the development of the Control of Agent-Based Systems (CoABS) Grid. The Grid is middleware that enables the integration of heterogeneous agent-based systems, object-based applications and legacy systems. The CoABS grid was used to develop the Critical Mission Data over Run-Time Infrastructure (CMDR) that allows dynamic discovery, integration and sharing of High Level Architecture (HLA) Run-Time Infrastructure (RTI) compliant simulation objects with legacy C4I systems and grid-aware software agents. The bridging of the CoABS grid and the HLA RTI using CMDR makes it possible to leverage the power of agent technology with the ability to tap into multiple C4I sources and simulation systems simultaneously. This synergy could lead to profound benefits in situation assessment and plan-execution monitoring using agents.

The key idea behind the capability as depicted in Figure 1 was to present the intelligent agents with real (notional) and simulated battlefield information, so that these agents can analyze both streams of data in order to understand and provide alerts when the notional battlefield information has changed as compared to the information contained in the data stream coming from the simulation [6]. Several types of agents were developed to check for positional deviations in the battlefield entities using the simulated data as ground truth, both using extrapolation and interpolation techniques. Extrapolation techniques -- to project the entity's real reported position -- were needed when the entity's reported time was less than the reporting time of that entity in the simulation in order to compare any changes to position at the same instance of time. Similarly, interpolation techniques were used when the entity's real reported position was greater than the reporting time of that entity in the simulation. An additional mass monitoring agent was responsible for detecting deviations in combat worth (number and value of assets) based on thresholds defined in the original plan in the simulation.

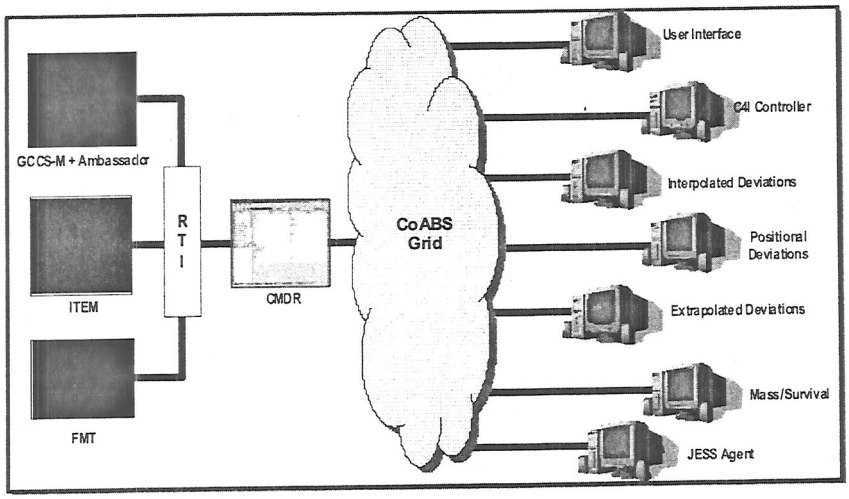

Figure 1: Federation of intelligent agents, C4I systems and simulations

This example demonstrates end-to-end system level coordination, facilitated by various middleware technologies such as HLA RTI, CMDR and CoABS grid. While the agents did not specifically communicate with each other in order to achieve some desired level of coordinated activity within the federation, the infrastructure does allow the agents to send and receive messages in order to enable their ability to communicate.

While the author's specifically describe agents to monitor plans, additional agents that are described for future development in the paper include plan-understanding agents. Through the development of appropriate plan-understanding agents, such agents would be able to determine critical events and relationships within a plan so that this information could be provided to the monitoring agents, which would then focus on monitoring the important aspects of a plan. The authors suggest that techniques such as natural language processing and sublanguage ontologies may be useful to extract key events from military operational orders or through leveraging XML-based techniques such the Battle Management Language [7].

2.4 Coalition interoperability

Coalition operations are characterized by data overload, concerns who to share information with, dealing with security issues and the need to overcome the stovepipe nature of systems -- all of which contribute to the challenges in achieving coalition interoperability. The goal of the Coalitions Agents Experiment (CoAX) was to demonstrate that coordination through multi-agent techniques could provide the mechanism to improve interoperability between systems. The CoAX [8] was an international effort between the Defense Advanced Research Projects Agency (DARPA), Defense Science Technology Office (DSTO), Defense Scientific Technical Laboratory (DSTL), and The Technical Cooperation Program (TTCP). The CoAX participants included the Department of Defense (DoD) laboratories, industry as well as academia, which were brought together with the purpose of demonstrating the value of multi-agent systems to support the ability to rapidly construct and maintain a coalition Command and Control structure. The experiment leveraged investment from the DARPA Control of Agent-based Systems (CoABS) program, specifically the CoABS grid framework.

In the fictional CoAX scenario two countries, Agadez and Gao, are in dispute over the territory of Binni, with Agadez becoming increasingly desperate over the territory. Because of this desperation, they launch a submarine-based missile strike against an Australian ship in the Red Sea. This strike in-

jures many on board and damages critical capabilities on the ship, including the Magnetic Anomaly Detection systems (MAD detectors). As the coalition becomes aware of the strike and subsequent damage, mobile medical monitoring agents are dispatched to the ship to collect injury reports contained in medical databases. From these reports, agents are able to work cooperatively to schedule the evacuation of those critically injured. A neutral country on the eastern coast of the Red Sea, Arabello, agrees to offer its ASW (Anti-Submarine Warfare) capabilities to help track down and neutralize the Agadez submarines. These ASW capabilities are rapidly integrated via agents with the coalition systems already in place, and expeditiously begin to provide contact reports on possible positions of the Agadez submarines in order to locate their positions.

In the context of the scenario, the goals of CoAX were to demonstrate that stovepipe barriers could be overcome across systems through techniques such as policy enforcement to bind the behavior of the agents in accordance with the coalition partner's requirements so that some form of trust could be achieved. For instance, services provided through the Knowledgeable Agent-Oriented System (KaOS) [9] assured that the agents would follow certain operating policies. The CoABS grid provided the infrastructure for the agents by providing registry, communication and logging services. There were dozens of agents that were developed and integrated using the CoABS grid in support of the scenario. In one particular instance within the scenario, agents must deconflict air plans, and this is accomplished through a multi-level coordination agent. This agent is built upon the notion of sharing plans that are represented hierarchically, and through the upward propagation of the conditions that hold true and the timing constraints between subplans, relationships between subplans at different levels of abstraction can be identified. Hence, agents can share these various subplans with each other at the appropriate level of detail in order to coordinate their activities.

Some of the goals of CoAX were to demonstrate that semantic web technologies could support interoperability between agents, an agent infrastructure (CoABS grid) could be used to rapidly integrate a diverse set of agents by providing support for communications and logging and that domain policies such as KaOS could lead to an ability to place a bound on the agents' behavior, which could eventually lead to trust relationships between the users and multi-agent system.

3 Transportation

3.1 Air Traffic Flow Management

Air traffic congestion is a world-wide problem leading to delays as well as subsequent monetary loss to the commercial flight industry. The ability to either manage the flow of traffic more efficiently or improve the infrastructure (e.g., adding more runways) to increase capacity are two available mechanisms to improve the efficiency of commercial aviation. Given the fact that increasing the capacity through improvements to the infrastructure is a costly and very time consuming solution, information technologies may be useful in helping to efficiently manage the traffic flow.

The Optimal Aircraft Sequencing using Intelligent Scheduling (OASIS) [10] is an intelligent agent-based system which is able to ingest data from Sydney Airport at the rate of 65 aircraft arriving in 3.5 hours. It is designed to assist the Flow Director (who is responsible for coordinating the landing sequences for aircraft) in arranging the sequencing of incoming aircraft through intelligent agents that represent the aircraft and supporting global agents. The aircraft agents are responsible for predicting aircraft trajectories, monitoring the actual movement versus predicted to understand discrepancies, and engaging in planning activities. The global agents consist of a coordinator agent, sequencer agent, trajectory manager agent, wind model agent and user interface agent. The coordinator agent is the central node for interacting with all of the other agents in the system. The sequencer agents uses search techniques based on the A^* search algorithm to determine the arrival sequence which minimizes delay and overall cost. The trajectory manager agent ensures that aircraft maneuvers do not violate statutory separation requirements. The wind model agent is responsible for computing and providing the fore-

cast winds based on the reported wind by the aircraft agents. Lastly, the user interface agent is the mechanism by which the flow director receives information from the system.

The planner agents accept a schedule time from the sequencing agent and decide how to execute that plan to meet that schedule, and these agents must monitor the actual movement against the plan and be able to predict the trajectory. If there are discrepancies, the aircraft agent must notify the global agents so that another sequence may be computed. The agents coordination is facilitated through communicate of their activities using the Procedural Reasoning System (PRS). The PRS system is an agent-based framework for agents to coordinate their activities with each other based on the Belief, Desires and Intention (BDI) model. The PRS framework is comprised of knowledge bases that contain agent beliefs as represented through first order logic, agent's goals or desires, a library of plans and a reasoning engine. Depending on the agent's beliefs and desires, the appropriate plans are triggered for execution, which in turn affect the environment.

The OASIS system is compared against COMPAS, MAESTRO and CTAS. The COMPAS system is developed for the German Civil Aviation Authority to provide flow control. It has been pointed out that the key difference between COMPAS and OASIS is that the former only handles single runway sequences and the aircraft's progress is not monitored. The MAESTRO system is developed for the French Civil Aviation Authority. Similarly to COMPAS, MAESTRO only considers a single runway, but does provide aircraft monitoring capabilities to provide a comparison of the progress against the plan. The CTAS is being developed by NASA-AMES for the US Federal Aviation Administration, and is described as more complex than either MAESTRO or COMPASS.

3.1 Other transportation and network management applications

Agents are deployed in this problem domain to cooperatively control and manage a distributed network. Agents can also handle failures and balance the flows in the network. Various approaches are developed in [11] to route packets in static and ad-hoc networks.
Distributed artificial intelligence (DAI) offers also suitable tools to deal with the hard transportation problems [12]. Fischer describes in [12] the modeling autonomous cooperating shipping companies system (Mars), which models cooperative order scheduling within a society of shipping companies. Three important instances for DAI techniques that proved useful in the transportation application, i.e., cooperation among the agents, task decomposition and task allocation, and decentralized planning are presented in this paper. Indeed, the problem tackled in this paper is a complex resource allocation problem. The auction mechanism is used for schedule optimization and for implementing dynamic replanning.

Coordination mechanisms are also used for distributed vehicle monitoring problems or traffic management. The objective in such problems is to maximize the throughput of cars through a grid of intersections or a network. Each intersection is equipped with a traffic light controlled by an agent. The agents associated to different intersections need to coordinate in order to deal with the traffic flows [13,14].
Another application concerns the Air fleet control or airspace deconfliction. A multi-agent approach for air fleet control is reported in [15]. In this application, the airspace used by the traffic controllers is divided into three-dimensional regions to guide the airplanes to their final destination. Each region can hold a limited number of airplanes. The multi-agent solution is used to guide the planes from one region to another along minimal-length routes while handling real-time data.

4 Healthcare – Monitoring Glucose levels

The ability to carefully monitor blood sugar levels is critical in the essential care of diabetes, which can lead to serious health problems and even death if not treated in a timely manner. A prototype applica-

tion called the Integrated Mobile Information System (IMIS) for diabetic healthcare is described in [6]. The IMIS demonstrates how multi-agent coordination can be leveraged to support critical collaboration between healthcare providers in the Swedish health care industry in support of monitoring diabetic patients.

Two types of problems are identified in [16] with regard to the proper care of diabetic patients, that of accessibility and interoperability. The accessibility problem is further categorized as physical service accessibility and information accessibility. The former implies that medical resources should be available to all healthcare providers while the latter implies that the appropriate information must be available from those resources. The key challenge identified in the diabetic healthcare system in one specific province, Blekinge, is information accessibility across the municipality and county council, each of which is responsible for various aspects of healthcare. For example, the municipality is responsible for the elderly, disabled or non-medical healthcare while the county council is responsible for medical and surgery related care. The various medical records cannot generally be shared between county council and municipality. The interoperability problem is one of overcoming the various forms in which the information is stored; in other words, there is a need for an information sharing standard to enable semantic interoperability.

The IMAS system is a multi-agent system that supports coordination between agents through the act of communication, in order to provide decision support to the human actors to support their collaborative activities. It was developed from the IMIS system and supports the collection and manipulation of healthcare information, and subsequent presentation of that information to the human actors based on their preferences in order to help them better collaborate. Three collaboration levels are identified in this specific problem domain: *Glucose management*, *Calendar arrangement* and *task delegation*. Glucose management involves providing the healthcare actors, perhaps through a centralized database, with the right information so that they can take the appropriate actions. Calendar arrangement includes the ability to schedule patient visits, while task delegation involves the ability of the healthcare providers or software agents to take certain actions (and the ability to report on the completion of those actions) based on the diagnosis.

There are four types of agents implemented in the IMAS system: *PatientAgent*, *ParentAgent, HospitalNurseAgent* and *SchoolNurseAgent.* These agents act on behalf of their user's preferences in order to better coordinate their activities through the communication of appropriate information (e.g., glucose levels). For example, in the scenario described in the paper [16], if a young boy (named Linus) has an elevated sugar level which is detected by the *PatientAgent* through readings from a database (which is updated by Linus via the IMAS patient control panel installed on his mobile device), then a message will be sent to the *ParentAgent* and *SchoolNurse Agent* to give Linus an insulin shot. When the alarm is received by the *ParentAgent,* this agent will search for additional information or charts to present to the mother, who will take the appropriate action. The IMAS system was implemented using the Java Agent Development Environment (JADE) and uses an ontology and uses the Foundation for Intelligent Physical Agents (FIPA) compliant messages for communication, and more specifically, the FIPA Agent Communication Language (ACL). The ACL is based on speech act theory, which is based on the fact that when one makes a statement, it is implied that one is doing something.

Future areas for investigation include techniques to ensure the privacy of the information so that only the right individuals or agents have access, and also the necessity to properly deal with the security of the information that is being transmitted so that it cannot be altered by third parties.

5. Communications Networks

5.1 Fault Diagnosis

As with any communication network, telecommunications networks are inherently distributed with expertise and data residing throughout various nodes of the network, hence, the management of the network requires access to, and analysis of, data that resides in various parts of the network. Due to the distributed nature of these networks, monitoring and diagnosing faults can benefit from a distributed approach as can be offered through multi-agent system techniques. The Distributed Intelligent Monitoring and Analysis System (DIMAS) system is described in [17], which builds upon the more centralized Intelligent Monitoring and Analysis System for monitoring and diagnosing faults. Specifically, the paper describes a distributed architecture based on a multi-agent approach for monitoring and detecting faults in cellular communication networks.

The problem which is described in the paper consists of monitoring and detecting faults across cellular base stations. Since multiple base stations may use the same frequency for communication due to the limited number of frequencies available, this has the effect of increasing the capacity of the network. Interference of frequencies across base stations is minimized by limiting the coverage and strength of the radio signal at each base station. However, in certain cases, natural phenomena such as winds or storms can cause the base station antenna to misalign. Such as misalignment can cause the frequencies to interfere, thereby affecting other nearby base stations by causing "noise" on the frequencies which overlap.

The DIMAS system is applied to monitoring this problem and diagnosing which base station is causing the frequency interference. The DIMAS system is comprised of a data store, data filter and expert system which is interfaced to a network simulator. The data store contains information obtained from the network, the data filter enables the processing of a large amount of information, and the expert system is the engine used to provide the diagnostic capabilities based on information obtained from the data store. The expert system in DIMAS is built on the AT&T C5 language, and its workflow processes through a series of steps: *monitor symptoms, reply to queries, do actions, analyze results, analyze evidence* and *decide next step*.

The authors also differentiate the DIMAS system with other related work. For example, they compare DIMAS to Distributed Big Brother (DBB). The DBB uses the Contract Net Protocol to distribute the monitoring task to Local Area network managers. However, within DIMAS, rather than offering tasks to potential bidders and then awarding the task to one bidder, the DIMAS agents are asked whether they can suggest possible actions to diagnose a problem. The suggested actions are then evaluated by the temporary owner of a problem and then contracted out to the agent that suggested the action. For example, each agent that detects a symptom of a problem may suggest further diagnostic tests, such as checking for changes to base station parameters or monitoring the signal strength of call attempts. However, for each action there is only one agent that can execute it. The decision that must be made by the coordinating agent is which action to choose, rather than which agent will execute that action. The subtle difference is that in DIMAS, the contracting is action-centric and not agent-centric, as a series of actions may be spread over multiple agents.

The Large-internetwork Observation and Diagnosis Expert System (LODES) is an expert system for diagnosing faults in a LAN that consist of multiple networks connected via routers. Each LAN has its own LODES system, which is used to detect remote faults based on local data. The primary difference between LODES and DIMAS is that the latter is a passive approach while the former is both passive as well as active (e.g., proactively sending test packets).

In DIMAS, the agents can only observe local symptoms and / or faults. Coordination between the agents is facilitated through the use of communication using KQML performatives such as *ask, tell* and *achieve* to gather non-local information. These performatives are used in various phases of the agent's process such as *assign-responsibility*, *gather-evidence* and *allocate-blame*. In the *assign-responsibility* phase, the problem is detected using the *ask* and *tell* performatives. The base station at which the

symptom is detected uses the *ask* performative to query the other base station agents about possible causes for increase in calls with poor transmission quality (frequency interference), while the other base stations use the *tell* performative to report back as to the fact that the cause may be due to a misaligned antenna or a frequency change. The affected base station agent then uses the *achieve* performative to ask the other base station agents to check for changes in frequency within their logs (*gather-evidence*). When this is ruled out, then more expensive requests are made such as checking for conditions which are indicative of a misaligned antenna. The *allocate-blame* phase assigns responsibility to the offending base station, and communicates its findings to other base stations.

5.2 Routing

A market-based approach to routing in telecommunication networks using a multi-agent system model is described in [18]. It is argued that a market-based approach requires no *a-priori* cooperation between agents and provides a robust mechanism for routing. Specifically, the approach relies on an adaptive price setting and inventory strategy based on bidding in order to reduce call blocking. The approach relies on several layers of agents, specifically, link agents that interact in link markets to purchase slices of bandwidth, path agents that interact in path markets to sell paths and on link markets to buy necessary links in order to offer paths on the path market, and lastly the call agents. Each of these agents buys or sells their commodity on the appropriate market.

The author compares the market based approach to routing with other approaches. For example, in a centralized scheme where traffic predictions are used to compute the paths there are issues with scalability particularly as the size of the network grows, the amount of information to monitor and process also increases and there is a single point of failure. Similarly, the authors argue that in optimal network flow algorithms there are performance issues in heavily loaded networks and unpredictable oscillations between solutions. Hence, the author suggests that a decentralized solution may be appropriate. However, the authors also describe a few challenges in decentralized solutions. For instance, they indicate that such solutions may have non-local effects on the entire network and it would be difficult to understand these effects due to the lack of complete information about the entire network. Complete information is infeasible since the network is dynamic and there is a delay in propagation and therefore the solution is prone to error and there may also be scaling issues.

As previously mentioned, the system architecture is comprised of link agents, path agents and call agents, each operating on one or several markets. A link agent and path agent is associated with each link and path, respectively, in the network. A call agent is associated with each source and destination pair in the network. Link and path markets are sealed bid double blind auctions in order to minimize any delays in the system, as callers will not want a delay in establishing a circuit in the network. The agents buy and sell the appropriate resource for which they are responsible, and depending on what they have sold and / or purchased, they use a simple function in order to either increase or decrease their buying / selling price. Similarly, path agents maintain an inventory that is profitable. Lastly, call agents trigger the auctions at the path and link agent level.

The authors present various results from their experiments. In the first experiment, they compare Vickery and First price auction strategies in their market-based approach to static routing on a 200 node network with link capacities for 200 channels. They show that they can do at least as well as static routing, but with the added advantage of achieving the performance level through an open architecture without the need to rely on static network policies. They demonstrate improved scalability through their approach, since the entire network topology need not be known in advance and, furthermore, using their approach are able to provide a quicker response as information is not needed about the broader network. The authors also provide results that demonstrate that about 61% of the time, the call was routed through the most efficient route and 46% of the time the call was routed through the least congested route. Research areas that are left for future work include the ability of the callers to

request different types of services, which might be routed differently through the network and the ability for market-based routing to support such activities.

6. E-Business

6.1 Supply Chain Management

A supply chain is described in [19] as "a network of suppliers, factories, warehouses, distribution centers and retailers through which raw materials are acquired, transformed, reproduced and delivered to the customer". A Supply Chain Management System (SCMS) manages the interaction of these processes in order to achieve some end goals.

The paper in [19] describes a SCMS based on the coordination of multi-agent systems, as represented through information agents and functional agents. The information agents are a source of information to the functional agents (for example, they provide a registry of other agents available in the system), while the functional agents are responsible for specific roles in the dynamic supply chain. The key notion behind these agents is that there is no assumption that a certain number of agents be available in the supply chain. The agents are assumed to enter and leave the system through a negotiation process based on the negotiation performatives that have been developed based on the FIPA ACL. For example, *accept-proposal*, *CFP*, *proposal*, *reject-proposal* and *terminate* are used for pair-wise negotiation between agents while a third performative, *bid*, has also been included to capture third party negotiations, for example, through an auctioneer.

The paper further describes the use of Colored Petri Nets (CPN) as a modeling tool for managing the negotiation process between agents. A simple example of the interaction between buyer and seller agents is described through the use of CPNs. The buyer agent initiates the negotiation process by sending a message indicating a *call-for-proposal* from the seller agent. The buyer agent then waits, while the seller agent receives the message for consideration. When the seller agent sends the reply (proposal, accept-proposal, reject-proposal or terminate) to the buyer agent, the buyer agent then enters the waiting or thinking / inactive state. If the thinking transition is triggered, the cycle repeats and it is up to the buyer agent to send the next round of message. The paper also describes the internal logic of the functional agents in making decisions by providing an example of an agent that represents a company's need for certain quantity of supplies in a given time period. This company is represented through an agent that sends a CFP message to other agents. The other agents evaluate the CFP and provide the quantity of supplies that can be provided, their cost and lead time. After receiving the other agents offers the initiating agent constructs a search tree for evaluating the constraints. If a solution can be found, then the initiating agent will accept one or more offers. If a solution cannot be found, either the initiating agent will relax its constraints, or ask the other agent to relax their constraints.

Several key issues are described for future research, including understanding the convergence behavior of the network and strategies for deciding how and when to relax the agent constraints.

6.1 Manufacturing Systems

Modern and big manufacturing depends heavily on computer systems. In fact, in many manufacturing applications the centralized software is not as effective as distributed networks. Indeed, to provide competitiveness in global markets, manufacturers must be able to implement, resize, design, reconfigure, respond to unanticipated changes, and maintain manufacturing facilities rapidly and inexpensively. These requirements are more easily satisfied by distributed small modules than by large monolithic systems. Moreover, small modules allow system survivability. Multi-agent systems offer a way to build production systems that are decentralized rather than centralized, emergent rather than planned, and concurrent rather than sequential.

Besides, the emergence of internet-based computing and communication to execute business processes has emerged as a key element in the supply chain domain. By using the internet, businesses gain more visibility across their network of trading partners, and it helps them to respond quickly to customer demands, resource shortages, etc. Min and Bjornsson present a method based on computer agents' technology for building virtual construction supply chains [20].

7. Emergency management and Disaster Relief

The emergence of new doctrine is enabling Security, Stabilization, Transition and Reconstruction (SSTR) operations to become a core U.S. military mission. These operations are now given equal priority to combat operations. The immediate goal in SSTR is to provide the local populace with security, restore essential services, and meet humanitarian needs. The long-term goal is to help develop indigenous capacity for securing and maintaining essential services. Therefore, many SSTR operations are best performed by indigenous groups with support from foreign agencies and professionals. Large scale disasters, however, are an example where military support can improve SSTR operations by providing a much needed boost to foreign governments and nongovernmental organizations which may already be under great stress to respond in a timely and effective manner. However, without the means to effectively coordinate the efforts between many diverse groups across the civil-military boundary during SSTR operations, basic assistance and relief operations may be severely impeded.

Many SSTR operational tasks are best performed by indigenous groups, with support from foreign or U.S. civilian professionals. Complex disasters are an example where military involvement and support for SSTR operations can provide significant value to foreign governments and non-governmental organizations (NGOs) which may already be under great stress to respond in a timely and effective manner. The command and control structure, resources and assets that the military can offer in such situations can shorten the response timeline. However, without the means to properly coordinate the efforts of such a large and diverse group which spans the civil-military boundary, basic assistance and relief operations may be severely impacted, leading to delays or waste in the overall response cycle.

There are many operational challenges in the ability to coordinate such a large group across the civil-military boundary during SSTR operations. Usually, in the civil sector, coordination is the result of voluntary effort; hence coordination by "directing" is rarely effective. Generally, relief agencies partly function within a framework of self-interest. In other words, they assist their beneficiaries in such a way that their good works are seen and valued by the donor community and the "profile" of their agency is enhanced. In many cases, farther down on the list is the goal of recognizing the contribution of others or admitting someone else can do the job better and therefore coordination is not necessarily an agencies first priority. With regard to coordination across the civil-military boundary there are also a number of challenges. The military tends to be a highly structured organization with a hierarchical command and control structure, while the civil sector in many cases is more loosely structured and tend to be less formal. These kinds of functional divisions can be confusing for each side involved in the coordination efforts. Furthermore, in the interest of national security, the military may be inclined to withhold certain information, while at the same time they may see it as the obligation of the civilian sector to provide as much information as possible. Hence, without adequate information exchange, coordination can become problematic.

With the signing of the Department of Defense Directive (DoDD) 3000.05, *Military Support for Security, Stabilization, Transition and Reconstruction (SSTR) Operations* into policy (U.S. DoD, 2007), SSTR operations have become a core U.S. military mission that the Department of Defense (DoD) must be prepared to conduct and support. These operations are now given equal priority to combat operations. Therefore, there is now a greater opportunity to overcome these fundamental cultural barriers and issues in order to move in the direction of proactive coordination. Therefore, one must understand the types of technologies that would be useful in order to provide decision support capabilities in order

to enable better coordination within the civilian sector, as well as across the civil-military boundary. Therefore, from a technology perspective, there are several areas that should be explored to enable more effective coordination in order to benefit those involved in SSTR operations.

Understanding emerging social networks and using the information that is provided through such networks is important for effective coordination. Social networks provide a visual cue as to the relationships between people or groups and can be used to identify "mutual" friends. This is particularly important when there is a requirement by one or more individual or groups to locate other individuals or groups that might be able to provide support in order to achieve a desired end result. Research and tools from the Social Network Analysis (SNA) community may allow users to understand who the experts are in the SSTR community and to whom and how they are linked. As a simple example, social maps that depict connectivity between users in the context of their discussion threads, and ability to filter the content within a social map based on specific keywords are likely to provide the foundation to enable the community to identify service providers or those that may offer similar services or capabilities. The ability to rate individuals within the social network may also be an important aspect in building trust within the community of users. This is particularly important during pre-disaster situations (i.e., prior to any crisis situation) so that some level of trust and common understanding can be achieved. Furthermore, pre-disaster interactions can help in the development of concept of operations or doctrine to provide guidance during real life situations by helping people or organizations form the bonds of working together. One of the challenges, however, will be to effectively visualize such a network or efficiently filter through the various dimensions of information contained in the social network.

There is also a lack of automated coordination tools to support the ability of people or groups to actively coordinate their activities. There are processes in place but most coordination is manual. The ability to manage tasks across the diverse actors involved in these operations has the opportunity to improve the overall efficiency of these operations. Due to the complexity of SSTR operations, multi-agent coordination techniques may provide benefits.

There is already ongoing research into developing information sharing environments to be used during SSTR operations within which automated coordination tools could be developed. These environments leverage content management systems and web 2.0 technologies to integrate various collaboration mechanisms such as weblogs, wiki's and forums. One of the key uses of such sites would be to enable service requestors to find service providers during crisis situations. As potential service providers are identified from the information sharing site's content, additional techniques could be developed and applied to transform and map the services that are being requested (again, from content posted on the site) into appropriate tasks. These tasks would be defined by appropriate users of the system, and would be defined such that when executed, would enable the services providers to meet the needs of the service requestors. The tasks may be hierarchical or peer-to-peer. For instance, hierarchical tasks may be appropriate for the military as protocols tend to be very structured and organized, while tasks for the civilian side may be peer-to-peer. However, the determination and selection of the appropriate task representation is best handled locally and there is no strict requirement that organizations use one representation or the other.

In order to automate task scheduling and assignment to a degree, the tasks and potential task performers might be organized into an assignment matrix, and the entries in the matrix would contain numerical values that represent the cost for doing each of the tasks across each of the performers (Figure 2). The costs may be monetary or may be based on a more complex cost function and it is up the participants to agree upon and choose the appropriate cost model. The computation of the cost functions across performers may vary, so some normalization techniques may need to be applied. Once these costs are determined, the goal of the matching algorithm would be to find a set of assignments between tasks and performers in the matrix that globally minimizes the total cost of doing the tasks across all of

Although MANET technology is advancing to enable connectivity between mobile users, there still may be circumstances in which users get disconnected (examples such as distance between users or the affects of the environment on signal propagation). In order to improve the overall success of the deployment of MANET, new approaches and techniques that enable users to communicate to the maximum extent possible utilizing whatever network bandwidth is available will be needed.

The concept of "NETwork-Aware" Coordination and Adaptation is a potential area worthy of exploration. In such an approach, the users or applications are aware of the state of the network, thereby allowing the applications to adapt in order to "work around" network constraints, while the network is aware of the state of the applications or mission needs in order to better handle traffic flows. Such cross-layer information exchange is important to enable a more robust communication strategy for the first responders in order to support their coordination activities. To the extent possible, coordination strategies also have to be robust against message loss and equipment failures.

A few of the research issues in network-aware coordination include defining measures for determining network congestion or other types of failures such as loss of connectivity within the network, in order to provide such measures and parameters to the application layer. The key challenges for the application layer include how to best utilize that information in order to adapt communication strategies (e.g., sharing images that are smaller in size, prioritizing certain information, or identifying certain nodes to act as communications relays). Such a feedback loop may be continuous, so that the network could support larger bandwidth exchanges as congestion is proactively alleviated in the network.

8. Tiered Systems

A key enabler of a sustainable military force is the notion of a tiered system. A tiered system is an integrated, multi-tier intelligence system encompassing space and air-based sensors linked to close-in and intrusive lower tiers [21]. The lower tiers (e.g., UAVs) are not only the critical source of intelligence; they can also serve as a key cueing device for other sensors. There is active research and exploration within the US DoD to understand the technical challenges in building tiered systems. It should be noted that tiered-system components such as UAVs or space-based assets are not only useful for ISR activities supporting more traditional combat operations, but may also enable effective SSTR operations.

Given the diversity of the assets, and the fact that coordination must be achieved both in the horizontal and vertical planes, and the environments in which the components of a tiered system will operate; it is not likely that a single coordination approach or even a family of coordination approaches will work well from a static perspective. It is more reasonable to expect that systems should learn which approaches work well and under which circumstances, and adapt appropriately.

It is commonly recognized that many different disciplines can contribute to a better understanding of the coordination function as a way to build and provide appropriate tools and adequate approaches for designing complex organizations and systems. Agents incorporated into these self-regulating entities represent "communities of concurrent processes which, in their interactions, strategies, and competition for resources, behave like whole ecologies". This paper has described a diverse set of application domains in which coordination of multi-agent systems has improved the behavior of the overall system. A common characteristic of these applications is the highly distributed nature of them.

9. Ambient Intelligence

The early developments in Ambient Intelligence took place at Philips in 1998. The Philips vision of Ambient Intelligence is: "people living easily in digital environments in which the electronics are sen-

sitive to people's needs, personalized to their requirements, anticipatory of their behavior and responsive to their presence".

The main purpose of Ambient Intelligence applications is to coordinate the services offered by small smart devices spread in a physical environment in order to have a global coherent intelligent behaviour of this environment. From a computing perspective, ambient intelligence refers to electronic environment that is sensitive and responsive to the presence of people. Ambient intelligence paradigm builds upon several computing areas and human-centric computer interaction design. The first area is **ubiquitous** or **pervasive** computing. Pervasive computing devices are very tiny (can be invisible) devices, either mobile or embedded in almost any type of object imaginable, including cars, tools, clothing and various consumer goods. All these devices communicate through increasingly interconnected networks (ad hoc networking capabilities). The second key area is **intelligent systems research.** Learning algorithms and pattern matchers, speech recognition and language translators, and gesture classification belong to this area of research. A third area is **context awareness**; research on this area lets us track and position objects of all types and represent objects' interactions with their environments. See 22 and 23 for more details about Ambient Intelligence applications.

10. Conclusion

Information technologies and information systems have become the cornerstone that shapes the present and the future of our society and its underlying affairs. They have penetrated almost every aspect of our lives. This omnipresence is stemming from significant advancements in computer systems, software, middleware, telecommunications infrastructure, etc. Also, enterprises, organizations and governments usually use a large variety of networked information systems, interconnected software, etc. From other side, military organisations as well as military coalitions share more information than before, such that decision making occurs at all levels within the chain of command. Indeed, current information technology gives organisations the opportunity to take advantage of all available information and coordinates the different actions. Consequently, coordination is becoming very important to consider and to take into account.

From IT research perspective the coordination paradigm has been usually studied by the community working in the area of Distributed Artificial Intelligence (DAI). This community unanimously endorses the idea that coordination is a fundamental paradigm that represents a challenging research area. A clear consensus has thus emerged around a forged, well-articulated, and strongly advocated common vision that coordination is a central issue to agent-based systems engineering research.

Emerging application domains such as *Ambient Intelligence*, *Grid Computing*, *Electronic Business*, *Semantic Web*, *Bioinformatics*, and *Computational Biology* will certainly determine future agent-oriented research and technologies. Over the next decade, it is thus expected that real-world problems will impose the emergence of truly open, fully scalable agent-oriented systems, spanning across different domains, and incorporating heterogeneous entities capable of reasoning, learning, and adopting adequate protocols of communication and interaction.

10. References

1. Liviu Panait and Sean Luke, "Cooperative Multi-Agent Learning: The State of the Art", In Autonomous Agents and Multi-Agent Systems, 11, 387-434, 2005.
2. Bedrouni, A., Mittu, R., Boukhtouta, A., Berger, J., Distributed Intelligent Systems: A coordination Perspective, Springer, New-York, USA, 2009

3. R. W. Beard, T.W. McLain, M.A. Goodrich, and E.P. Anderson, "Coordinated Target Assignment and Intercept for Unmanned Air Vehicles", In IEEE Transactions on Robotics and Automation, Vol. 18, No 6, pp. 911-922, December 2002.
4. Zhaotong Lian and M. Abhijit Deshmukh, "Performance prediction of an unmanned airborne vehicle multi-agent system", European Journal of Operational Research, No 172, pp. 680-695, 2006.
5. Goldman and S. Zilberstein, "Decentralized Control of Cooperative Systems: Categorization and Complexity Analysis". Journal of Artificial Intelligence Research, pp. 143-174, 22 (2004).
6. Mittu, R. Abramson, M. Walters. J., Improving Simulation Analysis through Interfaces to Command Control Communications Computer and Intelligence (C4I) Systems and Use of Intelligent Agents. Proceedings of the 2004 Spring Simulation Interoperability Workshop (SIW).
7. Carey, Scott., Kleiner, M., Heib, M., Brown, R., Standardizing Battle Management Language – Facilitating Coalition Interoperability. Proceedings of the 2002 Euro Simulation Interoperability Workshop (SIW),
8. Mittu, R., Segaria, F., Guleyupoglu, S., Barber, S., Graser, T., Ross, R., Supporting the Coalition Agents Experiment (CoAX) through the Technology Integration Experiment (TIE) Process", Proceedings of the 8th International Command and Control Research and Technology Symposium (ICCRTS), National Defense University, Washington, DC. 17-19 July 2003.
9. Feltovich, P., Bradshaw, J. M., Jeffers, R., & Uszok, A. (2003). Order and KAoS: Using policy to represent agent cultures. Proceedings of the AAMAS 03 Workshop on Humans and Multi-Agent Systems. Melbourne, Australia.
10. Magnus Ljungberg, Andrew Lucas, The OASIS Air Traffic Management System. Proceedings of the Second Pacific Rim International Conference on Artificial Intelligence (PRICAI '92).
11. Y.H. Chang, T. Ho, and L. Kaelbling, "Multi-agent learning in mobilized ad-hoc networks". Papers from the 2004 AAAI Fall Symposium, Technical Report FS-04-02-2004.
12. Fischer, "Cooperative Transportation Scheduling: An application Domain for DAI", Applied Artificial Intelligence, Volume 10, N1, pp. 1-34. 1996.
13. Nunes and E. Oliveira, "Learning from multiple sources", in AAMAS- Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems, 2004.
14. Dresner and P. Stone, "Multiagent traffic management; A reservation-based intersection control mechanism", In AAMAS- Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems, 2004.
15. Steeb R., Cammarata S., Hayes-Roth F., Thorndyke P., and Wesson R., "Distributed intelligence for air traffic control", In A. Bond and L. Gasser (eds.). Readings in Distributed Artificial Intelligence, Morgan Kaufmann Publishers, pp. 90-101, 1988.
16. Peng Zhang; Guohua Bai; Bengt Carlsson; Stefan J. Johansson pp. in proc. The International Conference on Information & Communication Technologies: from Theory to Applications - ICTTA'08, Damascus, Syria, 2008.
17. C. Leckie, R. Senjen, B. Ward and M. Zhao. "Communication and Coordination for Intelligent Fault Diagnosis Agents." In Proceedings of the Eighth IFIP/IEEE International Workshop for Distributed Systems Operations and Management (DSOM'97), Sydney, Australia, 21-23 October 1997, pp. 280-291.
18. M.A. Gibney, N.R. Jennings, N.J. Vriend, J.M. Griffiths. Intelligent Agents for Telecommunication Applications --- Proceedings of the Third International Workshop on Intelligent Agents for Telecommunication (IATA'99)
19. Chen, Y., Peng, Y, Finin, T., Labrou, Y., Cost, S., Chu, B., Yao, J., Sun, R., Wilhelm, B., A Negotiation-based Multi-agent System for Supply Chain Management". In Working Notes of the Agents '99 Workshop on Agents for Electronic Commerce and Managing the Internet-Enabled Supply Chain., Seattle, WA, April 1999."
20. Min J.U and Bjornsson H., Construction Supply Chain Visualization Through Web Services Integration, Technical Report, Civil and Environmental Engineering Dept., Stanford University, 2004.

21. Dahlburg, J., "Developing a Viable Approach for Effective Tiered Systems". NRL Memorandum Report 1001-07-9024. January 17, 2007
22. Malcolm McCullough, Digital Ground: Architecture, Pervasive Computing, and Environmental Knowing, MIT Press, 2005
23. Wang X., Dong J.S, Chin C., Hettiarachchi S.R. and Zhang D.. Semantic Space: A Semantic Web Infrastructure for Smart Spaces. IEEE Pervasive Computing, 3(3):32-39, July-September 2004