

AFRL-IF-RS-TR-2005-66
Final Technical Report
March 2005



COGNITIVE PERSONAL COORDINATION ASSISTANTS

Honeywell International Incorporated

Sponsored by
Defense Advanced Research Projects Agency
DARPA Order No. Q085

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

STINFO FINAL REPORT

This report has been reviewed by the Air Force Research Laboratory, Information Directorate, Public Affairs Office (IFOIPA) and is releasable to the National Technical Information Service (NTIS). At NTIS it will be releasable to the general public, including foreign nations.

AFRL-IF-RS-TR-2005-66 has been reviewed and is approved for publication

APPROVED: /s/

NANCY A. ROBERTS
Project Engineer

FOR THE DIRECTOR: /s/

JAMES A. COLLINS, Acting Chief
Advanced Computing Division
Information Directorate

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 074-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE MARCH 2005	3. REPORT TYPE AND DATES COVERED Final Jun 03 – Jul 04	
4. TITLE AND SUBTITLE COGNITIVE PERSONAL COORDINATION ASSISTANTS			5. FUNDING NUMBERS C - F30602-03-C-0010 PE - 61101E PR - COGN TA - A0 WU - 01	
6. AUTHOR(S) John Phelps, Thomas Wagner, Valerie Guralnik and Ryan VanRiper				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Honeywell International Incorporated Honeywell Laboratories 3660 Technology Drive Minneapolis Minnesota 55418			8. PERFORMING ORGANIZATION REPORT NUMBER N/A	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Advanced Research Projects Agency AFRL/IFT 3701 North Fairfax Drive Arlington Virginia 22203-1714			10. SPONSORING / MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TR-2005-66	
11. SUPPLEMENTARY NOTES AFRL Project Engineer: Nancy A. Roberts/IFT/(315) 330-3566/ Nancy.Roberts@rl.af.mil				
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.				12b. DISTRIBUTION CODE
13. ABSTRACT (Maximum 200 Words) COORDINATES are coordination managers from human teams. Each human in a team is paired with a COORDINATOR coordination manager which is running on a mobile computing device. COORDINATORS provide decision support to teams by reasoning about who should be doing what, when, with what resource, in support of which other team, and so forth. COORDINATORS respond to the dynamics of the environment by (re)coordinating to determine the right course of action for the current circumstances. We implemented COORDINATORS using wireless PDAs and proprietary human/asset location tracking technologies. This report discusses the concept of COORDINATORS, the motivation for them, the underlying agent architecture, implementation issues, the first response exercises we used to test them, and ideas for future work.				
14. SUBJECT TERMS Coordination, Agents, Decision-Support, Real-Time			15. NUMBER OF PAGES 29	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

Table of Contents

1. Introduction.....	1
2. COORDINATORS Technical Exploration.....	1
2.1 Introduction.....	1
2.2. Motivationand Business Concepts.....	4
2.3 Agents, Architecture, and Implementation.....	7
2.4 Domain Expertise.....	9
2.5 Coordination and First Response Exercises.....	10
2.6 COORDINATORS Extensions, Limitations, and FutureWork.....	14
2.6.1 Strategic COORDINATORS Problems and Models.....	14
2.6.2 Advanced Tactical COORDINATORS Problems and Models..	17
2.6.3 Learning for COORDINATORS.....	21
References.....	24

LIST OF FIGURES

1.	COORDINATORS help first responders coordinate joint action by reasoning about tasks and interactions.	2
2.	A network of COORDINATORS handling task coordination between responders. .	3
3.	A single COORDINATOR running on an Off-The-Shelf wireless PDA.	3
4.	A COORDINATORS first responder value tree.	4
5.	Ansoff's Product/Market expansion grid.	6
6.	A Single TÆMS Agent / the core of a COORDINATOR.	7
7.	Overview of the MAS architecture of a network of COORDINATORS.	8
8.	Locator tag that tracks humans/assets for the COORDINATORS system.	9
9.	Creating a Rescue Mission by instantiating a mission template.	10
10.	Our development system enables displaying both the "Mobile" device side interfaces and centralized command on a single workstation.	11
11.	A single slice of larger coordination episode: Tasks and schedules reflect <i>Team₂</i> 's support of <i>Team₃</i>	12
12.	Domain work for the exercise is carried out with props.	13
13.	After the exercise, first response stand-ins are debriefed and shown the optimal course of action and the choices they made.	15
14.	Initial TÆMS strategic view.	16
15.	Initial TÆMS tactical control view.	17
16.	A proposed Task List interface.	18
17.	A proposed Change Task interface.	19
18.	Modified TÆMS tactical control view.	20
19.	A Propsed New Task Alert dialog.	21

1. Introduction

This document summarizes the DARPA program creation effort of COORDINATORs. It also includes a description of technical exploration done for COORDINATORs, including prototype deployment, coordination protocol development, coordination problem theory, and future work.

2. COORDINATORs Technical Exploration

2.1. Introduction

This section describes the technical exploration work that was done to support the DARPA program creation task. Dr. Wagner's team at Honeywell developed ideas and a demonstration of TÆMS agent technology for the First Responder domain that provided important inputs for the creation of the DARPA program. This section covers the following: COORDINATORs application motivation, TÆMS agents and architecture for COORDINATORs, COORDINATORs First Responder domain description, description of fielded exercises of COORDINATORs, thoughts on using COORDINATORs for strategic and tactical applications, learning, and unaddressed research issues.

For the technical exploration task, COORDINATORs were designed to support the coordination of first responders such as fire fighters or police. They provide decision support for first response teams and the incident commander by reasoning about mission structures, resource limitations, time considerations, and *interactions* between the missions of *different teams* to decide who should be doing what tasks and when so as to get the best overall result. COORDINATORs provide global team activity optimization – helping the teams to respond to the dynamics of the environment and to act in concert, supporting one another, as appropriate for the current circumstances. When the situation changes, the COORDINATORs communicate, evaluate the implications of change, and potentially decide (or suggest, depending on their role) on a new course of action for the teams.

There are several characteristics of this problem instance that make it a hard problem:

The situation is dynamic – it is not known with any detail at the time of the 911 call what sort of state the site or victims will be in when response teams arrive. Thus the agents must coordinate and decide which operations to perform in real-time. This is especially true when fire is involved; in an unmitigated average office fire, gas temperature inside the burning, enclosed space can easily reach 1200 degrees Fahrenheit in less four minutes[8].

Agents must make quantified / value decisions – different tasks have different values and require different amounts of time and labor resources. It may be critical to provide water supply support to suppress fire spread until victims are discovered during a search, at which point, priorities require adjustment.

Coordination is dynamic – the operations being performed by the first responder teams interact and the occurrence of the interactions are also not known *a priori*. For instance, until victims are found, it is not known whether ventilation in a hallway will be required.

Deadlines are present – a fire suppression team will need to put out a fire in one area within a deadline in order for a rescue operation to be able to effectively complete their evacuation operation. Deadlines require the agents to reason about end-to-end processes and to coordinate with other agents to optimize their activities.

Tasks are interdependent – tasks interact in two different ways: 1) over shared resources in a spatial/temporal fashion, 2) multiple tasks must be performed to accomplish a goal, e.g., a fire has not been met with a satisfactory response until all the people threatened by it have been evacuated, and it has been extinguished in the most effective manner possible (though in TÆMS this generally pertains to degrees of satisfaction rather than a boolean or binary value).




Figure 1. COORDINATORs help first responders coordinate joint action by reasoning about tasks and interactions.

The underpinnings of COORDINATORs are TÆMS agents [2, 17, 18, 7] equipped with a new coordination module derived from the coordination keys [17] technology. This means that each distributed COORDINATOR is able to reason about complex mission task structures and communicate with other coordinators to determine who should be supporting whom, when, in order to save the most lives, make the best use of assets or resources, reduce risk to the response teams, and so forth. An application-centric view of an individual COORDINATOR is shown in Figure 1. A network of COORDINATORs is shown in Figure 2.

Implementationally, COORDINATORs have been constructed using off-the-shelf wireless PDAs and are currently being ported to more specialized wearable computing devices.




Figure 2. A network of COORDINATORs handling task coordination between responders.

COORDINATORs also leverage a Honeywell-proprietary asset location technology to track the physical location of first response teams, victims, and important resources such as a wall cutting saw or a multi-story portable ladder. A screen snapshot of a PDA-based coordinator is shown in Figure 3. The screen is showing the location of the teams with its owner team's schedule arranged across the top.

A simple demonstration of COORDINATORs for first responders is implemented and functioning and has been experimented with using staged first response exercises. However, this project and the work described here is only the potential starting point for COORDINATORs and technology that supports human activity coordination. The demonstration explored only a small subset of the larger problem space of *cognitive* COORDINATORs that learn to improve, reason about organizational structures when decision making, reason about change in the environment, and exhibit other advanced reasoning capabilities.



Figure 3. A single COORDINATOR running on an Off-The-Shelf wireless PDA.




Figure 4. A COORDINATORs first responder value tree.

Note that herein we use the term “first responder” to mean personnel ranging from fire fighters to emergency medical teams. For the details of this project, however, we have focused primarily on the needs of the fire fighters and the incident commander because we were able to get domain expertise in that area.

We now move on to discuss the first response domain, the value proposition for COORDINATORs and some technology development issues. We then provide architectural and technical details of COORDINATORs and discuss human-based first response exercises using COORDINATORs.

2.2. Motivation and Business Concepts

COORDINATORs are for large-scale first response team coordination. To frame the problem space, imagine a large-scale crisis event such as a terrorist attack, at a large facility such as a university campus or a petrochemical plant, with multiple concurrent incidents, and multiple organizations or teams responding. In situations such as this, effective response requires coordination between the first response teams. They must act both in concert, supporting each others’ efforts (and attempting not to hinder one another), and individually, carrying out their own different assignments. In this environment, the situation is changing in real-time and there is often not very much *a priori* information available. This means that the teams and the incident commander must form and adapt their plans online as the situation unfolds. Currently, this process is carried out using walkie-talkies and the teams rely heavily on the incident commander to provide high-level coordination. The problem is that reasoning about who should be doing what, and when, in such complex, distributed situations is very difficult for humans. (Consider the number of hours per week one spends simply scheduling meetings.) The problem gets worse as team sizes grow. Add-in that the situation is highly dynamic and that change requires timely evaluation and response. Then factor-in the crisis element –

flames, explosions, and human lives at stake. Human decision making in these situations translates into high levels of cognitive load, coarse approximation in reasoning, and, from an evaluation standpoint, suboptimal coordination. This is a situation in which decision support technologies can help make better decisions, faster, with a greater attention to detail (finer grain of coordination), and with a (near) optimal utilization of teams and resources. The goal of COORDINATORs is to enable the human responders and the incident commander to focus on the human-hard problems and to off-load coordination reasoning on to the automated coordination managers.

This concept has been well received by experts in fire, security, and first response, as well as by other customers for military or industrial applications. Keep in mind, however, that the technical vision motivating COORDINATORs is *long term*. This concept is still very far away from something that could be used by first responders in the field. Consider the device level issues alone – heat, moisture, *steam*, darkness, device interaction in a loud setting while wearing gloves and carrying equipment, device power, ad-hoc networking/connectivity, are just a few of the areas that must come together before this concept is fully viable. These issues *are* being addressed by other research communities but interaction with experts and our own product divisions has made it clear that this vision is still several years away from deployment.

Figure 4 shows the value tree for the COORDINATORs concept. While the different dimensions of value can generally be regarded as the result of efficiency improvements and having adjustable optimization criteria, such structures help clarify the origin of potential value of an investment. The tree depicts that the value propositions for COORDINATORs stem from improvements they provide in the quality, speed, cost, and risk for emergency response. Quality is improved because COORDINATORs enable First Responders to consider a larger space of response options with the repercussions of choices in time spelled out. The combinatorics of such calculations quickly overwhelm humans without COORDINATORs. Considering the larger space of options leads to a higher probability of choosing options that save more lives and protect civilian and First Responder equipment and other property.

COORDINATORs also run faster than human coordination due firstly to the compact encodings of task options. The increased speed is due to the fact that computer communication and analysis of response options is much faster than human communication and analysis – this becomes more pronounced as the team size scales. The computer encoding of task models in TÆMS, including potentially numerous alternatives, leads to less time required for analysis, coordination, and thus lead to an overall faster response. The reduction in response times and ability to consider the impact of response choices, enables COORDINATORs to produce a more efficient response force, thus reducing response costs.

Finally, the value tree depicts that COORDINATORs lower the risk to First Responders,




Figure 5. Ansoff's Product/Market expansion grid.

civilians, and assets. This is due to two factors. Perhaps most importantly, COORDINATORs provide more effective means of exploring a larger number of alternative courses of action and their implications, as mentioned above. This alone reduces risk inherent in the reactive and error-prone human coordination without COORDINATORs. Secondly, COORDINATORs rely on the TÆMS representational framework, which can encode expectations about task outcomes as distributions. COORDINATORs can use that probabilistic information when considering alternatives to manage uncertainty in the selected TÆMS quality, cost, or duration performance dimensions.

Another important tool when developing a new technology concept like COORDINATORs is Ansoff's product-market-expansion grid, shown in Figure 5. The general idea is that the current or present market/product mix is in the upper-left hand side and expansion moves from that point to one of the other spaces in the grid. The move that most businesses will balk at is the move from the upper-left to the lower-right. This means creating a new product and trying to sell it in a new market. This is regarded as a risky proposition because they are not working from a position of strength. As a technical person, developing a concept that moves a company into that lower-right box means that investment may be very difficult to obtain. On the other hand, a move from the upper-left to the upper-right is a move that business development people endorse. This means taking a current product and developing a new market for it. This is an instance of expansion from a strength. However, this expansion movement is generally of little use to those developing technical concepts because, by definition, we are generally creating new products and new ideas. The one expansion avenue open to work like COORDINATORs is to keep the company in its existing market (movement from upper-left to lower-left) but to develop a new product for that market. Whether or not COORDINATORs actually makes the move to the lower-left is a matter of some debate – partially because the entire first response / homeland defense market is in a state of flux. What is helpful when developing a concept like COORDINATORs is to hit on multi-purpose use – something that is an enabler on a day-to-day basis as well as something that has value in first response and crisis situations.




Figure 6. A Single TÆMS Agent / the core of a COORDINATOR.

2.3. Agents, Architecture, and Implementation

COORDINATORs are based-on TÆMS agents [2, 17, 18, 7]. The name TÆMS stands for *Task Analysis Environment Modeling and Simulation* though in its present day usage TÆMS is best regarded as a hierarchical modeling language used to represent complex task networks and interactions between different agents. TÆMS-based components, such as the Design-to-Criteria (DTC) agent scheduler [16, 19, 12] then reason about these complex task networks to decide on a course of action for the agent. TÆMS enables TÆMS-based technologies to be used in many different domains because, as with any modeling language, it provides a layer of abstraction away from the domain details. Figure 6 shows a single TÆMS agent as constructed for COORDINATORs. From a high-level, each TÆMS agent contains sophisticated control problem solving modules that reason about tasks, task interactions, interactions that span agents, time deadlines, resource constraints, and so forth to decide who should be doing what, and when, so as to optimize the activities of a group of distributed agents. The details of how the technologies operate are beyond the scope of the present discussion, though more information can be found in detailed discussions of TÆMS [2, 7], DTC agent scheduling [16, 19, 12], GPGP agent coordination [2, 1, 6], and a similar approach to team coordination [17].

From the larger Multiagent System (MAS) view, each COORDINATOR has a single TÆMS agent to which it is linked via wireless 802.11b network, as shown in Figure 7. In the future, we envision the TÆMS agents themselves running on portable computing devices but the processing requirements of a TÆMS agent are too great for current PDAs and wearable computers (shrinking the TÆMS agent requirements has not been explored to conserve resources). The current COORDINATOR system architecture has each wireless PDA connecting to a TÆMS agent back-end – the PDA agent is essentially an interface stub that communicates as needed with the TÆMS agents. Each TÆMS agent in turn may communicate with other TÆMS agents to coordinate the activities of different teams (exchange local information, negotiate over task interactions, determine who should be doing what, and when). When the system is running in simulation mode (as it is generally

as we have few buildings to actually burn for experimental purposes), agent communication is routed through the Multiagent Survivability Simulator (MASS) [15] simulation environment (as shown in the figure). Each TÆMS agent also communicates with the Honeywell Lab’s building Enterprise Building Integrator (EBI) server. The EBI server is a commercial proprietary technology for controlling building systems and for tracking assets within a facility. For instance, managers within the building wear asset location tags so that they can be located as needed. In the COORDINATORs system, first responders are tagged so they can be tracked through the facility and so the teams can see visually where other teams are located and track each other’s movements using the PDAs and digitized maps. The incident commander COORDINATOR (another TÆMS agent supports it) can also track team movements using the asset location technology. The addition of situational icons, e.g., locations of known fires, is planned for future implementation. Obviously, the information gathered from the tracking system could also be fed into other automated reasoning modules in a more comprehensive COORDINATOR, e.g., one that provides path planning or intelligent evacuation support using the information.

An asset location tag is shown in Figure 8. The tag uniquely identifies a person or an asset by sending a signal to receiver units which are distributed throughout the building. Location of the party in question is determined by proximity to a given receiver. In many modern office buildings / facilities, such tagging is commonplace and becoming more so. Note that there are many other location technologies that might be used in systems like COORDINATORs, including GPS and cellular-based triangulation (both of which generally perform poorly indoors).

COORDINATORs are implemented in a combination of Java and C++. The reasoning components, e.g., the DTC agent scheduling module, rely on the Java Agent Framework (JAF) and the MASS simulation environment [15] to provide inter-component state access and inter-agent message transport. The JAF infrastructure includes a Java representation




Figure 7. Overview of the MAS architecture of a network of COORDINATORs.

of TÆMS. DTC, the major C++ component, uses its own internal TÆMS representation. As discussed, the TÆMS agents themselves run on Linux-based workstations while the PDA interactions are managed by Java-based “agent-lets” or stubs. The advantage of this design and these platform choices is that we do not need the PDAs to run the rest of the system – the PDA-side Java applications can simply be run under linux and connected to using a physical network, as shown in Figure 10. This facilitates testing, experimentation, and debugging.

The user interfaces developed for COORDINATORs attempt to strike a good balance between research/development cost and functionality. Recall that we regard these devices as demonstration vehicles – for actual deployment a different device is needed and user interaction will most likely *not* be based on stylus manipulation. Figure 9 shows the steps necessary to create a rescue mission. The data entered via pull- down menu and check box selection serves to fill-in the details of an existing mission template. We are experimenting with alternative PDA interfaces (a discussion of which is given in a later section), speech recognition, heads-up displays, and wearable computers as a way to facilitate more natural environment-embedded interaction.

2.4. Domain Expertise

Our data on first response and large-scale events comes partly from first hand interviews (by another effort at Honeywell Labs) with fire marshals and first responders. Other sources of information include reports on the 9/11 attack, the Oklahoma bombing, and other similar events. Still other information is obtained from published first response plans or procedure documentation of different first response agencies. While domain expertise has helped to shape and motivate this work, there is room for more grounding as the work matures. A likely next step is to perform the demonstrations with and for practicing first response personnel. Previous interactions with first responders and those versed in first response lead to the recognition that the incident commander’s expertise must be leveraged and harnessed, not by-passed. That the COORDINATORs must support a centralized commander probing deep into the activities of each team and potentially (re)tasking them as appropriate. The COORDINATOR network would then respond to these changes and adapt all the other team’s actions to dovetail with the commander’s directives.



Figure 8. Locator tag that tracks humans/assets for the COORDINATORs system.




Figure 9. Creating a Rescue Mission by instantiating a mission template.

2.5. Coordination and First Response Exercises

While a detailed example of first responder coordination in COORDINATORs is beyond the scope of the present discussion (one can be found in [21]), the critical conceptual point is that coordination requires reasoning about tasks and interactions. That is, understanding how one team’s tasks interact with those of other teams. These interactions require cohesive choice (everyone selecting the “right” way to perform tasks) and temporal sequencing (everyone must do their actions at the “right times”). From a high-level, coordination is about deciding *who should be doing what*, and *when*, so as to get the best overall result for the *current* circumstances. The idea behind COORDINATORs is that they continuously evaluate change in the situation or in the mission and decide how the teams should adapt and respond. COORDINATORs are a technology intended for *online* use which means they must respond in “soft” real-time and cannot rely on algorithms requiring exponential search. A small example of a coordination episode is shown in Figure 11. At this point in time (in a larger scenario), *Team₂*’s COORDINATOR has negotiated with *Team₃*’s COORDINATOR for *Team₂* to support *Team₃*’s evacuation efforts by providing secondary lighting for a darkened stairwell. Later in the same scenario, due to new deadlines imposed by the potential for explosions elsewhere in the facility, the COORDINATORs must renegotiate and consider alternatives – *Team₂* winds up supporting *Team₃* by providing a net for window evacuation instead of lighting for




Figure 10. Our development system enables displaying both the “Mobile” device side interfaces and centralized command on a single workstation.

stairwell evacuation.

In terms of evaluation, arguably the most important overall evaluation question for COORDINATORs is whether they improve the performance of first responders. In a perfect world with unlimited resources, one might design a set of experiments in which first responders engage in a series of first response episodes both with and without COORDINATORs providing support. In each case, one would like to measure specific metrics like number of lives saved, number of assets saved, time required to perform the mission tasks, number of responders necessary to address the situation, amount of risk incurred by the responders and the civilians, etc. In this perfect world, one would have buildings to burn and the ability to recreate, verbatim, scenarios so that the measurement and comparison could be one-to-one.

We elected to use an approach that somewhat realistically simulated the sorts of coordination problems first responders might encounter. To evaluate COORDINATORs from an application view, rather than simply evaluating the performance of the underlying technology (e.g., time required for coordination), we staged first response exercises and had human performers take the role of first responders. Note that the lessons learned from this process are anecdotal but are also more meaningful as an early viability test of the concept.

In the exercises there are four teams and an incident commander (IC). The scenario is set in a petrochemical plant though the plant is mapped back onto the Honeywell Lab’s building. During the exercise, responders must move around the building, perform




Figure 11. A single slice of larger coordination episode: Tasks and schedules reflect *Team₂*'s support of *Team₃*.

situation assessment tasks, respond to the situations they discover, and coordinate to rescue civilians. The scenario is setup in such a way that teams must coordinate in order to rescue the civilians. Failure to do so results in (simulated) loss of life – a metric that can be tabulated.

To assess the benefits of having COORDINATORs, we first deploy the teams on the first response exercise using walkie-talkies for communication (they are also equipped with stop-watches and building maps to make the simulation more complete). After the walkie-talkie exercise, during which loss of (prop) life is recorded, the teams are rotated and the scenario run again, this time with COORDINATORs providing automated support.

In doing this exercise, we rapidly discovered the degree to which humans are overwhelmed when faced with lots of temporal and task related data that is in a state of constant change. The initial plan was to have each participant take the role of incident commander – the individual who generally handles coordination in the walkie-talkie exercise. Not only was the IC task too difficult for most participants, it was too difficult for most of the research team members. In practice, only someone who had memorized the flow of events in the exercise could help the teams to rescue all the civilians. We resorted to this model in order to get human performers through the walkie-talkie exercise at all.

Thus the participants (with varying degrees of domain expertise) generally took the role of first response teams. At the start of the scenario, the teams are deployed by the IC and given situation assessment tasks. In enacting the scenario, at this point teams move throughout the building and go to assigned zones (generally conference rooms). To simulate the situation assessment task, we created a series of props representing the situation. For




Figure 12. Domain work for the exercise is carried out with props.

instance, a first response team might find fire props, debris props, and a civilian prop pinned by a girder prop. This would indicate that a civilian was trapped and that the fire needed to be put out and the debris cleared before the girder could be cut away. Cutting the girder also requires some other team (generally) to fetch a power saw from the simulated truck. In the exercise, props are reinforced by staging data sheets that describe the situation textually and explicitly cull out resource needs and potential temporal issues (e.g., “you must evacuate these civilians before the adjacent wall collapses at time $T=40$ ”).

Because fielded first responders must coordinate while carrying out domain tasks, we also require our first response stand-ins to carry out simulated domain tasks. In general, this translates into putting props into one another and moving them physically throughout the building. Figure 12 illustrates the process of putting out a small fire. To extinguish the fire, it goes into a *fire extinguishment box* and the box must then be carried to a staging area on a specific floor of the building. Similarly, evacuation of an injured civilian requires that the civilian prop be put into the gurney prop box, a box that must be fetched from the staging area, and then the gurney box must be put into a stairwell box (if that is the exit route chosen) and the stairwell box carried to the staging area.

Dynamics are introduced into the environment using secondary envelopes on which is printed a time at which they are to be opened. Thus teams may coordinate, decide on a course of action, then open an envelope and discover that the situation has changed (e.g., a ceiling fell-in) and then they must recoordinate to adapt to the new situation.

As one might guess from the description, human performers generally fared poorly during this exercise. Only with an expert IC who knew the complete scenario *a priori* and had figured out exactly who should be supporting whom, and when, could get both the teams

and the cardboard civilians out of the facility in time. What is more interesting is that the stress incurred by the human performers during the exercise was pronounced and observable even to the non-expert. Trying to battle one's props while processing all the cross chatter on the walkie-talkie and interact with the IC proved to be a difficult task even without the heat, smoke, sound, and inherent danger of a crisis situation. Few performers were able to coordinate properly. Few were able to evaluate their mission structures properly. Not once did a guest team make it through the scenario with the optimal course of action chosen.

In contrast to the walkie-talkie scenario, the run with COORDINATORs handling the activity coordination is almost boring – despite the scenario being run at a faster clock rate. In the COORDINATOR scenario, the teams perform situation assessment and describe their situation to the COORDINATORs. The COORDINATORs then handle all of the exchange of local information, the analysis, and the formation of commitments. Teams are then informed of what they should be doing, when, who will be supporting them, and so forth.

After both exercises, the participants are then debriefed and shown a simplified Gantt chart, Figure 13, of the major coordination points and support needs of the different teams. While the evidence gathered during these exercises is anecdotal, the reaction of our visitors, some with first response and military domain expertise, has served to reinforce our belief that this line of work is valuable. In practice, the “fog of war” caused by flames, screaming, smoke, etc., makes a set of tasks that humans have difficulty with under normal circumstances nearly impossible. Information exchange and coordination analysis should be off-loaded from the humans to automated assistants that are better equipped to reason precisely and respond in a (near) optimal and timely fashion.

2.6. COORDINATORs Extensions, Limitations, and Future Work

We have presented COORDINATORs, discussed the motivation behind them, and identified important COORDINATORs development value concepts. We have also described the underlying TÆMS agent technology and described the anecdotal experimentation with COORDINATORs. In this section we discuss extensions and limitations to COORDINATORs models, interfaces, and reasoning capabilities. We conclude by summarizing some promising directions for future work.

2.6.1. Strategic COORDINATORs Problems and Models

Although our implementation of COORDINATORs was focused on tactical response coordination, the underlying technology may be used with similarly powerful effect on strategic coordination problems. Modeling emergency response scenarios strategically is a way to abstract away aspects of the emergency response problem to provide a tractable view of the problem at the level of incident commander coordinating with other incident




Figure 13. After the exercise, first response stand-ins are debriefed and shown the optimal course of action and the choices they made.

commanders or above. TÆMS supports analysis at different levels of abstraction naturally through its hierarchical structure and the uniformity of structural attributes: characteristics, accumulation functions, and non-local effects.

Uniformity of TÆMS structural attributes enables the COORDINATORs system modeler to digest lower problem complexity into higher-level abstractions, e.g., aggregating performance bounds into a “high-level” TÆMS method. Conversely methods with expected (or directed) characteristics, can be decomposed into more detailed task structures prescriptively, according to templates created before the start of a mission or that result from online planning. If within a high-level simulation scenario, a lower level of problem abstraction is required to meet certain quality, cost, and deadline constraints, those constraints can be input directly into the local problem solver’s task structure analysis constraints.

We now develop an example that shows how a strategic TÆMS view can map to a tactical TÆMS view, and how a Generalized Partial Global Planning (GPGP) coordination mechanism operates over the tactical views. Although we don’t develop a case for coordination at the strategic level in this example, the same mechanisms apply to the strategic level due to the recursive structure of TÆMS and the uniformity of the task




Figure 14. Initial TÆMS strategic view.

structure attributes. In our example, we suppose a fire breaks out in a ten story hotel, Hotel 13. This event, through a dispatch call, results in the instantiation of the task structure depicted in Figure 14. Our response model is derived from the response model of the Boston Fire Department [3]. We base our responding units on the Boston Fire Department’s Third District composition [4]. The Third District is one of eleven districts in Boston, together operating many different types of response units from ladders and engines to emergency medical response teams to structural collapse units [4].

If a TÆMS strategic model were to be developed in the field for pre-incident emergency response planning, values for a default response plan could be based on historical data and the judgement of the response personnel. In reality, the response model for an in-building fire includes deploying 3 Engine Companies, 2 Ladder Companies, 1 Rescue Company, and a District Chief [3]. We model a subset of the actual response. Realism is not done away with, since elements of the response team may arrive asynchronously in some situations, especially when supporting other districts or in large, spreading fire or another sort of extremely taxing emergency. In the default response plan for the part of the Boston Fire Department’s Third District hypothetical emergency response team that we have modeled,




Figure 15. Initial TÆMS tactical control view.

special tasks for Hotel 13 could be included, as incident response plans often do. Here, for simplicity's sake, we show a more generic response plan. The plan includes, for a fire emergency, generally, responding to the fire in the building and supplying medical care to individuals. The general tasks are decomposed to task structure abstractions embodied in methods characterized through historical data and human expertise.

Critically, after the decomposition of tasks, which may be accomplished through formal decentralization policies [22], the composite task structure attribute values from a tactical level may propagate back up to the strategic level through a mutual commitment algorithm [13] or delegation, e.g., Engine-8 company will report the results of all of the locally responding companies from the Third District. Tasks are distributed naturally at the strategic view according to the resources they require in this domain.

In this scenario many of the tasks required to respond to a single event need to be performed in sequence. Several sets of tasks cannot be performed simultaneously because they involve the same spatial or temporal areas as prerequisite tasks. For instance, rescue operations cannot be performed in areas where the entire structure is in jeopardy from a spreading fire. In contrast, many other tasks can be performed asynchronously, including containing the fire, helping evacuated victims, and connecting auxiliary water sources.

2.6.2. Advanced Tactical COORDINATORs Problems and Models

While strategic task modelling and analysis abstracts away from the details of the requisite tasks in an emergency situation, basing its analysis on abstract information and predesignated values, tactical task modelling and analysis for COORDINATORs is about coordinating workflow "in the trenches." Tactical tasks are at a finer grain size, and, for fire and rescue personnel might include extinguishing fire in a section of a building,



Figure 16. A proposed Task List interface.

containing a fire to a building section, searching for victims, or helping evacuated victims.

We continue with our fire and rescue scenario by breaking out some of the tactical tasks for two teams and then showing how these tasks are coordinated using extensions to our existing PDAs mockups that stress some of the unique capabilities of the TÆMS agent technology that underpins COORDINATORs.

From the strategic view given in Figure 14, it is clear that the incident commander has decided to take an offensive approach. The initial thrust will be spent trying to extinguish and contain the fire inside the building in anticipation of the search and rescue teams. There will be three waves, based on the schedules obtained from the available personnel. The ladder will arrive first, shortly after that an engine will arrive with a search and rescue team. Figure 15 is the tactical view of Ladder-1 company.

The task list interface shown in Figure 16 shows a first responder’s current task list. This is an extension of the schedule view that we implemented in our COORDINATORs prototype. In the idea presented, each task description contains a label. Above, tasks with labels **Suppress Window Fire** and **Check Adjoining Hallway** are listed. The arrows to the right of the tasks provide a way of navigating through the different tasks.

The **Done** button provides a way to indicate that a task is complete. There is no way to indicate early completion of a task in the current implementation. Further, any interaction with tasks is done in a separate dialog, rather than in the same mode as the task list display.

The **Change** button would bring up a **Change Task** dialog, shown in Figure 17. The -1:27 and -3:22 fields in the task list indicate how much time is remaining for each task. There is also an importance field for each task, which is marked **HIGH** for **Suppress Window**




Figure 17. A proposed Change Task interface.

Fire and MED, indicating medium relative importance, for **Check Adjoining Hallway**. Although we relegated the function of changing the priority of a task to the **Change Task** dialog to emphasize that this change might involve a “challenge-and-response” interaction, the priority could also easily be peremptorily adjusted in this mode. That approach ignores important issues about how those priority changes affect the priority of other tasks in the COORDINATORs network and how those changes are arbitrated.

A success probability for each task is also given. Representing uncertainty, although difficult for humans to estimate, can enable the underlying TÆMS agent system to dynamically explore contingency plans for tasks with low probability of success. A key research question for future work is how to deal with the uncertainty associated with uncertainty estimates. Machine learning techniques seem especially fitted to address this issue.

The **New** button would allow a responder to add new tasks to the network. The **Clear** button clears all tasks, indicating a problem encountered that could be annotated with radio communication or understood through new tasks coming into the network. A key idea that this technology leverages is that responders would be able multitask/multiplex communication with several teams, often implicitly through task network template instantiation, as is shown in Figure 18.

The **Change Task** activity coordination screen is shown in Figure 17. This screen would allow an emergency responder to modify a task that he or she was currently tasked with. The responder could request a deadline adjustment, or could change the probability of task success, or even challenge its importance. Those changes could then either be accepted through the **Accept** button or cancelled through the **Cancel** button. What happens at that point is that the new details are forwarded to other responder’s agents, and new




Figure 18. Modified TÆMS tactical control view.

commitments are negotiated if possible. If the agents could not resolve a commitment revision, an audio warning could be issued and a dialog between the emergency response personnel involved could ensue to resolve the task change contention.

To develop the tactical scenario a bit more deeply, we demonstrate how the TÆMS agent network could manage task revision dynamically. During a check of the second floor room that contained the window fire, a member of the Ladder-1 company discovers that there are people trapped behind a small fire in the room. The Ladder-1 company member then creates a task to rescue people, called **Rescue People in Second Floor Room** in Figure 18. This task instantiates with options that can be quickly tuned to the given situation. This one includes preconditions for rescue, including extinguishing the flame in the room and venting the hallway.

The new tasking could cause a nearby company member who is nearly done extinguishing a window fire, to receive a **New Task Alert**. An example of what this could look like with advanced interface options is shown in Figure 19.

The **New Task Alert** describes a task to vent the hallway, with a deadline in one minute, a preestimated probability of success, and high importance. The **Accept** button cannot be selected, however, because accepting this task would cause the responder, with high probability, to miss the deadline on the preexisting **Check Adjoining Hallway** task. In this case, a default timeout could cause the **New Task Alert** to close if there is no response. We have discussed adding a “validity time” counter to the arriving tasks in our next iteration of the COORDINATORS system.

Alternatively, the responder could click the **Explore Feasibility Options** button to adjust or cancel the conflicting tasks. In a feasibility exploration mode, various pending




Figure 19. A Proposed New Task Alert dialog.

commitment requests and scheduled tasks could be looked at in combination. Various attributes of the tasks in a particular combination of tasks or across combinations of tasks could be adjusted and their effects considered. The responder could then select the combination that provided the highest utility according to the most up-to-date estimates of local task performance. If time was scarce, the COORDINATORs system could perform a heuristic search to find the best combination based on high-level criteria provided by the responder.

2.6.3. Learning for COORDINATORs

Machine Learning can be applied in several areas of TAEMS agents. Several general approaches to learning in agent systems have been explored [5, 11, 14], and some exploratory work on learning task structures in agent systems has also been done[9]. One area includes learning default task assignment rules. For instance, if nearly every time a fire department gets a call Ladder-1 arrives first, then Engine-8 arrives, a default task assignment could be made so that Engine-8 could plan ahead to perform supporting tasks more efficiently. For instance, if a secondary source of water is usually required, the Engine-8 company could ready themselves to provide the secondary water source. The company could be immediately routed to the source rather than requiring a coordination interaction with Ladder-1. As another example, let's say that a member of Ladder-1's company, Sue, often asks a co-member Bob to assist her in rescue operations and Sue frequently negotiates over commitments, Bob would not want to learn a default commitment.

Another area where machine learning can be useful is to learn the rules for selection of an agent with which to negotiate a commitment. Imagine that you don't have a fully specified task structure but instead are given TAEMS tasks with *enables* specified as a need, e.g.,

Need Safety Net. If the agents then have to find other agents that can satisfy their requests through a discovery protocol – this might be especially useful in large, dynamic disaster situations. A T&EMS agent could learn a preference for requests of certain items. Those preferences could be propagated throughout the organization. For instance, it could be learned to prefer not to ask for slack resources from a normally highly constrained nearby district or department. Similarly, organization roles may be learned [10].

Hard problems are abundant in COORDINATORS. At the most basic level, distributed activity coordination from partial views in a real-time setting is an extremely difficult problem – particularly given the degree of complexity needed to represent mission tasks and their interactions. This is the reason for the “nearly optimal” caveats in this paper. Most techniques that we have developed to date are approximate but perform well on average, e.g., the keys coordination mechanism [17]. Another hard issue that we just began to investigate in COORDINATORS is how to provide a centralized command-and-control interface to a network of COORDINATORS that are inherently distributed – issues of how much information and control is shared between COORDINATORS must be resolved as well as how authority relations are encoded. Through our interaction with domain experts it became clear that we needed to support and leverage the incident commander, not replace him/her. A related issue is how to support mixed modes of decision making – COORDINATORS need to be able to (learn to) make decisions for their teams in some circumstances and in others need to consult their team members (or the IC) directly. This might involve advanced situation assessment or reasoning about what the team is currently doing, the importance of the decision at hand, the existence of other options, the response time required for the decision, and so forth. Another issue hereto unaddressed is the classic question of “where do the models come from?” Our assumption is that we can create a library of templates and then instantiate them at run time – possibly asking the IC or the responders to adapt the missions to the current situation. This is time consuming and will require the right interfaces, right mission editing tools, etc., in order to make it feasible (our current tools could be improved). Even if these issues can be resolved, they are still predicated on the assumption that missions can be generalized and they repeat (the hypothesis that this is the case is based on the current day existence of *response plans* and standardized procedures). Another area of future work involves organizational structures – reasoning about decision making procedures and following proper organizational structure, and decision making protocols, for both COORDINATORS and humans.

In the immediate future, we plan to explore some of the underlying technical issues in COORDINATORS in greater depth. For instance, the coordination algorithm used here has only been evaluated in the small and is known to have some unaddressed issues, e.g., considering load balancing when tasking. An obvious extension to that work will be to evaluate COORDINATORS with more complex tasks and larger teams.

References

- [1] Keith Decker and Jinjiang Li. Coordinated hospital patient scheduling. In *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS98)*, pages 104–111, 1998.
- [2] Keith S. Decker. *Environment Centered Analysis and Design of Coordination Mechanisms*. PhD thesis, University of Massachusetts, 1995.
- [3] Boston Fire Department. Alarm responses. In <http://www.ci.boston.ma.us/bfd/divisions/alarm-responses.htm>. Internet URL, January 2003.
- [4] Boston Fire Department. Firehouse locations. In <http://www.ci.boston.ma.us/bfd/divisions/locations.htm>. Internet URL, January 2003.
- [5] Cora B. Excelente-Toledo and N. R. Jennings. Coalition, cryptography, and stability: Mechanisms for coalition formation in task oriented domains. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002)*, pages 1106–1113, 2002.
- [6] Victor Lesser, Keith Decker, Thomas Wagner, Norman Carver, Alan Garvey, Daniel Neiman, and Nagendra Prasad. Evolution of the GPGP Domain-Independent Coordination Framework. Computer Science Technical Report TR-98-05, University of Massachusetts at Amherst, January 1998.
- [7] Victor Lesser, Bryan Horling, and et al. The TÆMS whitepaper / evolving specification. <http://mas.cs.umass.edu/research/taems/white>, July 2004.
- [8] National Fire Protection Association. *Fire Protection Handbook*, 18th edition, 1997.
- [9] Utgoff Paul, Jensen David, and Lesser Victor. Inferring task structure from data. In *International Conference on Machine Learning*, January 2000.
- [10] M. V. Nagendra Prasad, Victor Lesser, and Susan Lander. Learning organizational roles in a heterogeneous multiagent system. In Sandip Sen, editor, *Working Notes for the AAAI Symposium on Adaptation, Co-evolution and Learning in Multiagent Systems*, pages 72–77, Stanford University, CA, 1996.
- [11] M. V. Nagendra Prasad and Victor R. Lesser. The use of meta-level information in learning situation-specific coordination. In *IJCAI (1)*, pages 640–646, 1997.
- [12] Anita Raja, Victor Lesser, and Thomas Wagner. Toward Robust Agent Control in Open Environments. In *Proceedings of the Fourth International Conference on Autonomous Agents (Agents2000)*, 2000.
- [13] Tuomas W. Sandholm. Distributed rational decision making. In Gerhard Weiss, editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, pages 202–258. MIT Press, 1999.
- [14] Peter Stone and Manuela M. Veloso. Multiagent systems: A survey from a machine learning perspective. *Autonomous Robots*, 8(3):345–383, 2000.

- [15] Regis Vincent, Bryan Horling, and Victor Lesser. An agent infrastructure to evaluate multi-agent systems: The java agent framework and multi-agent system simulator. In Thomas Wagner and Omer Rana, editors, *Infrastructure for Agents, Multi-Agent Systems, and Scalable Multi-Agent Systems, Lecture Notes in AI*, pages 102–127. Springer, 2001.
- [16] Thomas Wagner, Alan Garvey, and Victor Lesser. Criteria-Directed Heuristic Task Scheduling. *International Journal of Approximate Reasoning, Special Issue on Scheduling*, 19(1-2):91–118, 1998. A version also available as UMASS CS TR-97-59.
- [17] Thomas Wagner, Valerie Guralnik, and John Phelps. A key-based coordination algorithm for dynamic readiness and repair service coordination. In *Proceedings of the 2nd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS2003)*, 2003., 2003. Nominated for most novel application award.
- [18] Thomas Wagner, Valerie Guralnik, and John Phelps. Software Agents: Enabling Dynamic Supply Chain Management for a Build to Order Product Line. *International Journal of Electronic Commerce Research and Applications, Special issue on Software Agents for Business Automation*, 2(2):114–132, 2003.
- [19] Thomas Wagner and Victor Lesser. Design-to-Criteria Scheduling: Real-Time Agent Control. In Wagner/Rana, editor, *Infrastructure for Agents, Multi-Agent Systems, and Scalable Multi-Agent Systems*, LNCS. Springer-Verlag, 2001. Also appears in the 2000 AAAI Spring Symposium on Real-Time Systems and a version is available as University of Massachusetts Computer Science Technical Report TR-99-58.
- [20] Thomas Wagner, John Phelps, Valerie Guralnik, and Ryan VanRiper. An application view of coordinators: Coordination managers for first responders. In *Proceedings of the Sixteenth Innovative Applications of Artificial Intelligence Conference (IAAI04)*, 2004.
- [21] Thomas Wagner, John Phelps, Valerie Guralnik, and Ryan VanRiper. Coordinators - coordination managers for first responders. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS04)*, 2004.
- [22] Ping Xuan and Victor Lesser. Multi-agent policies: from centralized ones to decentralized ones. In B. Drabble, editor, *Proceedings of the first international joint conference on Autonomous agents and multiagent systems.*, pages 1098–1105. ACM Press, 2002.