

An Application View of COORDINATORS Coordination Managers for First Responders*

Thomas Wagner

John Phelps

Valerie Guralnik

Ryan VanRiper

Honeywell Laboratories
John.Phelps@honeywell.com

Abstract

COORDINATORS are coordination managers for fielded first responders. Each first response team is paired with a COORDINATOR coordination manager which is running on a mobile computing device. COORDINATORS provide decision support to first response 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. COORDINATORS have been implemented using wireless PDAs and proprietary first responder location tracking technologies. This paper describes COORDINATORS, the motivation for them, the underlying agent architecture, implementation issues, and first response exercises.

Track: Emerging Applications

Topic Area: Software Agents

Introduction

COORDINATORS are coordination managers for fielded first responders such as fire fighters. 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, 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

* Effort sponsored by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL) under agreement number F30602-03-C-0010 and by Honeywell International under project numbers I10155-AB and I10133-AC. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency (DARPA), the Air Force Research Laboratories, the U.S. Government or Honeywell International.
Copyright © 2004, American Association for Artificial Intelligence (www.aaai.org). All rights reserved.

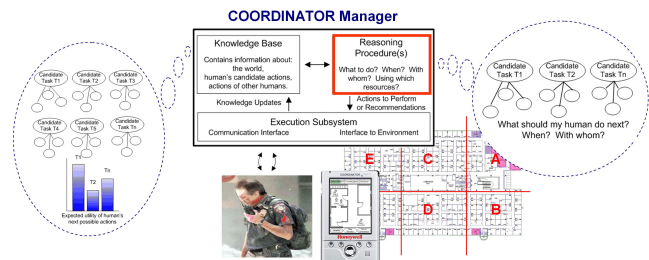


Figure 1: COORDINATORS Help First Responders Coordinate Joint Action By Reasoning About Tasks and Interactions

of change, and potentially decide (or suggest, depending on their role) on a new course of action for the teams.

The underpinnings of COORDINATORS are TÆMS agents (Decker 1995; Wagner, Guralnik, & Phelps 2003a; 2003b; Lesser, Horling, & et al) equipped with a new coordination module derived from the coordination keys (Wagner, Guralnik, & Phelps 2003a) 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. COORDINATORS also leverage a Honeywell proprietary asset location technology (akin to RF id tags) 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.

COORDINATORS are implemented and functioning and have 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.

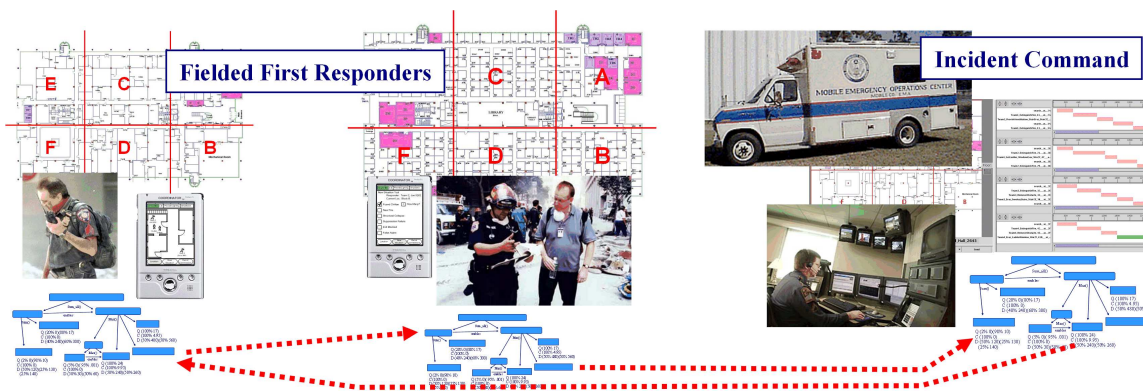


Figure 2: A Network of COORDINATORS Handling Task Coordination Between Responders

DARPA (Defense Advanced Projects Research Agency) is exploring the possibility of creating a larger effort in coordination which would lead to *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.

Note that throughout this paper 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.

In this paper we discuss the first response domain, the motivation for COORDINATORS, and a few underlying business development issues. We then provide architectural and technical details of COORDINATORS and discuss human-based first response exercises using COORDINATORS.

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 a large scale situation is very difficult for humans. (Consider the number of hours per week one spends simply scheduling meetings.) 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 a factual 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



Figure 3: A Single COORDINATOR Running on an Off-The-Shelf Wireless PDA

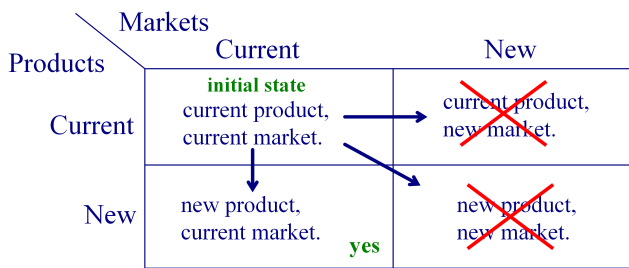


Figure 4: Ansoff's Product/Market Expansion Grid

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.

Another important tool when developing a new technology concept like COORDINATORs is Ansoff's product-market-expansion grid, shown in Figure 4. 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 business people 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. Whether COORDINATORs for first response will achieve product status in the distant future is far from predictable.

Agents, Architecture, and Implementation

COORDINATORs are based-on TÆMS agents (Decker 1995; Wagner, Guralnik, & Phelps 2003a; 2003b; Lesser, Horling, & et al). 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 DTC agent scheduler (Wagner, Garvey, & Lesser 1998; Wagner & Lesser 2001; Raja, Lesser, & Wagner 2000) 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 5 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 this paper, though more information can be found in papers on TÆMS (Decker 1995; Lesser, Horling, & et al), DTC agent scheduling (Wagner, Garvey, & Lesser 1998; Wagner & Lesser 2001; Raja, Lesser, & Wagner 2000), GPGP agent coordination (Decker 1995; Decker & Li 1998; Lesser *et al.* 1998), in a different COORDINATORs paper (Wagner *et al.* 2004), and in the aircraft service team coordination application (Wagner, Guralnik, & Phelps 2003a).

From the larger MAS system view, each COORDINATOR has a single TÆMS agent to which it is linked via wireless 802.11b network, as shown in Figure 6. 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 commu-

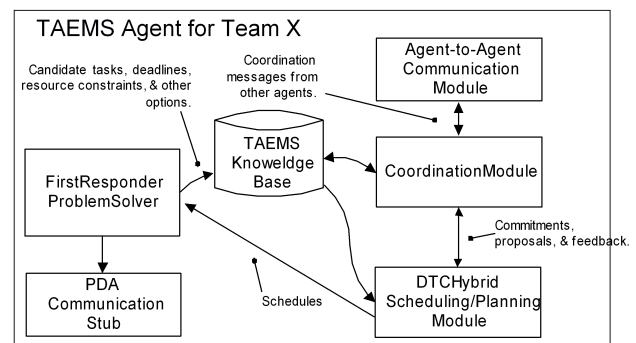


Figure 5: A Single TÆMS Agent / the Core of a COORDINATOR

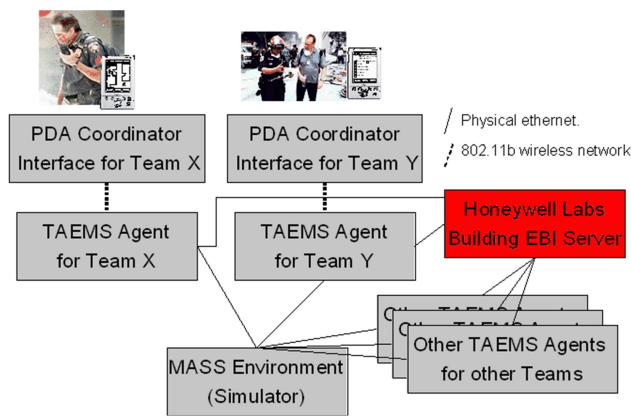


Figure 6: Overview of the MAS Architecture of a Network of COORDINATORS

nication is routed through the MASS (Vincent, Horling, & Lesser 2001) simulation environment (as shown in the figure). Each TÆMS agent also communicates with the Honeywell Lab's building 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. These ideas have not been explored in any meaningful way to date. An asset location tag is shown in Figure 7. 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 (Vincent, Horling, & Lesser 2001) to provide inter-component glue/messaging/execution and inter-agent message transport. The JAF infrastructure includes a Java representation 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



Figure 7: Locator Tag Tracks Humans/Assets & Their Locations

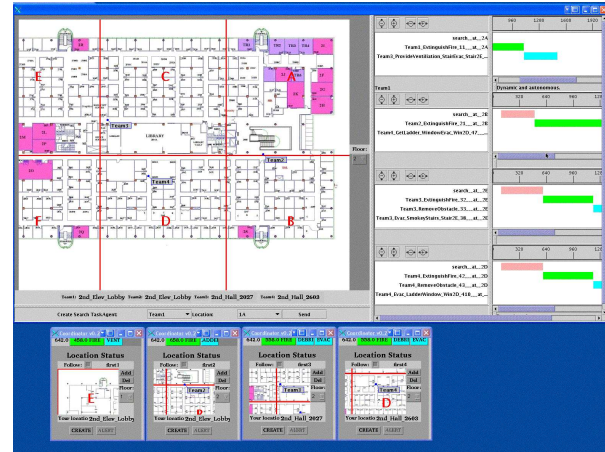


Figure 9: Java Enables Displaying Both "Mobile" Device Side Interfaces and Centralized Command on a Single Workstation

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 9. 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 8 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 speech recognition, heads-up displays, and wearable computers as a way to facilitate more natural environment-embedded interaction.

Coordination and First Response Exercises

While a detailed example of first responder coordination in COORDINATORS is beyond the scope of this paper (one can be found in (Wagner *et al.* 2004)), the critical conceptual point is that coordination requires reasoning about tasks and interactions. That is, understanding how one team's tasks

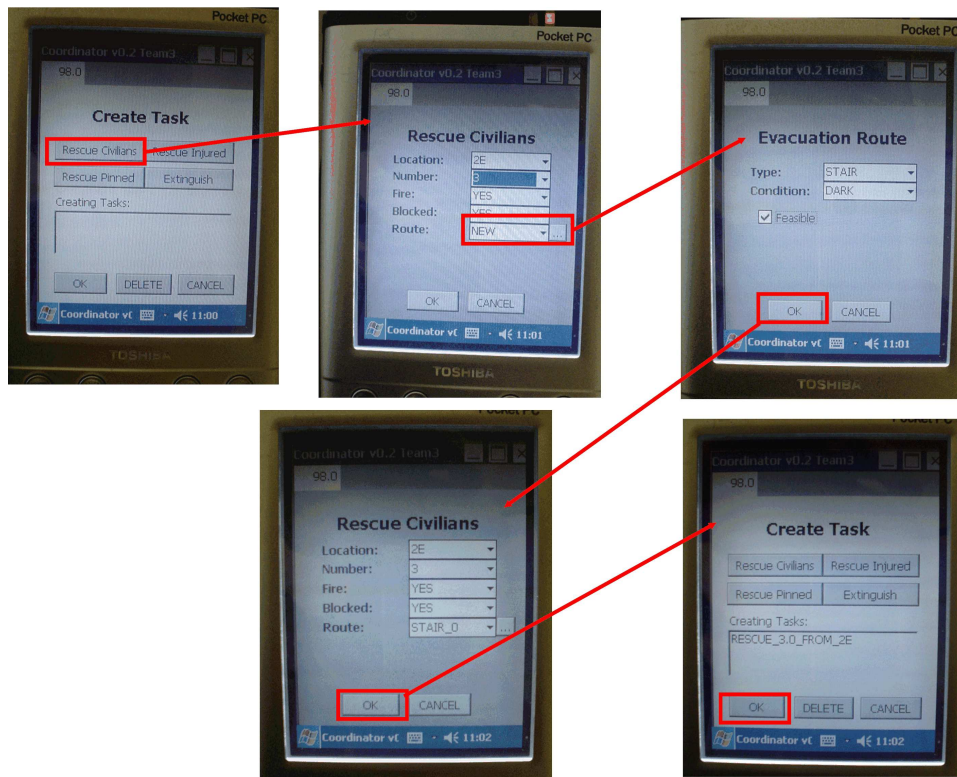


Figure 8: Creating a Rescue Mission By Instantiating a Mission Template

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 10. 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 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 COORDI-

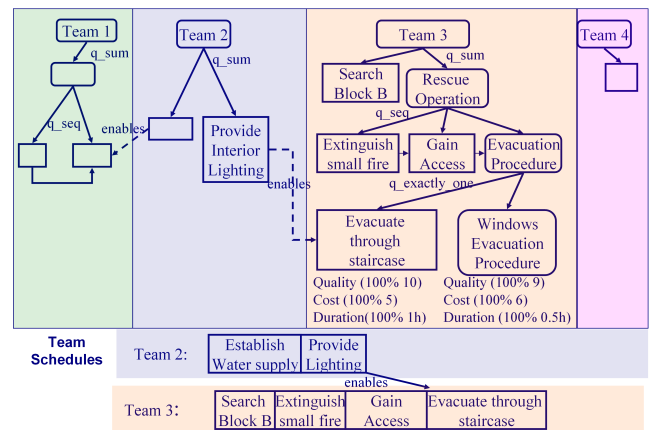


Figure 10: A Single Slice of Larger Coordination Episode: Tasks and Schedules Reflect *Team₂*’s Support of *Team₃*

NATORS 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.

Not living in a perfect world, we elected to use a somewhat more economic approach. To evaluate COORDINA-

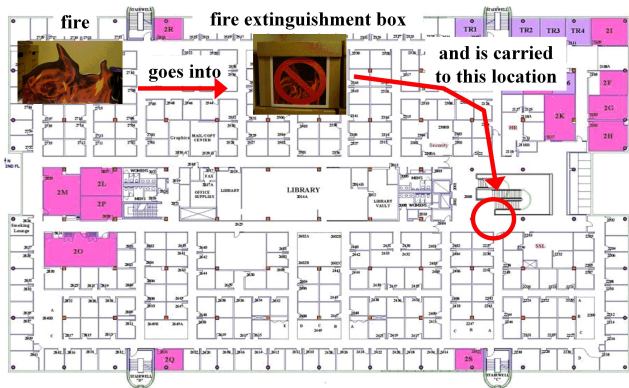


Figure 11: Domain Work is Carried Out With Props

TORs 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 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 host VIPs and to have a VIP 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 the VIPs, 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 VIPs and visitors (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 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 11 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 re-coordinate 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. Notable among our VIPs was a Honeywell VP who processed the temporal data without hand drawn Gantt charts and who carried the props with great vigor while barking commands into his walkie-talkie. (Confidence in management rose a fraction during this episode.)

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




Figure 12: After the Exercise, First Response Stand-ins are Debriefed by Being Shown the Optimal Course of Action and the Choices They Made

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 VIPs are then debriefed and shown a simplified Gantt chart, Figure 12, 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.

Unaddressed Research Issues, Limitations, and Future Work

In this paper, we have presented COORDINATORS, discussed the motivation behind them, and identified a few important business development concepts. We have also described the underlying TÆMS agent technology and described the anecdotal experimentation with COORDINATORS.

COORDINATORS relate to several areas of prior art. The electronic elves (Chalupsky *et al.* 2002) project underscored the value of using a portable computing device as a front end for a more complex, agent-based, support system. The use of GPS tracking in e-elves also affirmed our interest in leveraging Honeywell asset tracking technologies in COORDINATORS. Our work differs from e-elves in that the reasoning and information exchange is all distributed – in the-

ory no agent has a complete view of the activities of the other agents. COORDINATORS also differs in the type of task/temporal analysis required to coordinate complex mission structures. In a sense, COORDINATORS are more narrowly scoped than e-elves but require deeper analysis technology. The use of agents and portable computing devices to disseminate information to, and collect information from, aircraft repair teams (Shehory *et al.* 1999) also contributed intellectually to COORDINATORS. Interfacing with the building systems and displaying digital maps, along with team location, information is a concept that relates naturally with this prior work. Another area is the work done in using TÆMS agents to support aspects of distributed collaborative design processes (Decker & Lesser 1995). Though there are differences, e.g., grainsize and response time requirements, the underlying technical problems are probably members of the same class or related classes.

While this paper has described the COORDINATORS project, we have not dealt overmuch with research issues (more technical details appear in (Wagner *et al.* 2004)). Hard problems are abundant in COORDINATORS. At the most basic level, distributed activity coordination from partial views in a (soft) real-time setting is an extremely difficult problem – particularly given the degree of complexity needed to represent mission tasks and their interactions. This the reason for the “nearly optimal” terminology in this paper. Most techniques that we have developed to date are approximate, e.g., the keys coordination mechanism (Wagner, Guralnik, & Phelps 2003a). Another hard issue is how to provide a centralized command-and-control interface to a network of COORDINATORS that are inherently distributed. 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 future, we plan to explore some of the underlying

technical issues in COORDINATORS in greater depth. For instance, the coordination algorithm used here has not been evaluated in the small and it is known to have some unaddressed issues, e.g., considering load balancing when tasking. The next steps for this work will also be partially influenced by DARPA interests and Honeywell interests. This work was partly funded by DARPA's IPTO office as an exploratory project in coordination. Their interest is in creating greatly enhanced *cognitive* COORDINATORS that not only solve all the core problems well, but that learn and adapt over time, reason about organizational structure and military decision making when coordinating, and are scalable by learning to control their own problem solving (learning with whom to coordinate, learning which coordination decisions are important, etc.). From the cognitive COORDINATOR perspective, the COORDINATORS discussed here should be viewed as an early prototype of things to come and what is possible in this problem space.

Acknowledgments

Victor Lesser's group at UMASS is an industrial affiliate of Honeywell and a subcontractor on the COORDINATORS project. TÆMS has a long history and has involved many researchers, including: Victor Lesser, Keith Decker, Alan Garvey, Tom Wagner, Bryan Horling, Regis Vincent, Ping Xuan, Shelley XQ. Zhang, Anita Raja, Roger Mailler, and Norman Carver. We would also like to acknowledge the efforts of the other Honeywell Lab's project team who performed first hand interviews with fire marshals and first responders. These include Tom Plocher, Walt Heimerdinger, Tony Faltesek, and Michelle Raymond. This work was partially supported by DARPA's IPTO office, through AFRL, and by Honeywell International. We would like to thank Ron Brachman, Barbara Yoon, Zach Lemnios, John Beane, Mike Lynch, and Andrew Berezowski for their support.

References

- Chalupsky, H.; Gil, Y.; Knoblock, C.; Lerman, K.; Oh, J.; Pynadath, D.; Russ, T.; and Tambe, M. 2002. Electric Elves: Agent Technology for Supporting Human Organizations. *AI Magazine* Summer.
- Decker, K. S., and Lesser, V. R. 1995. Coordination assistance for mixed human and computational agent systems. In *Proceedings of Concurrent Engineering 95*, 337–348. McLean, VA: Concurrent Technologies Corp. Also available as UMASS CS TR-95-31.
- Decker, K., and Li, J. 1998. Coordinated hospital patient scheduling. In *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS98)*, 104–111.
- Decker, K. S. 1995. *Environment Centered Analysis and Design of Coordination Mechanisms*. Ph.D. Dissertation, University of Massachusetts.
- Lesser, V.; Decker, K.; Wagner, T.; Carver, N.; Garvey, A.; Neiman, D.; and Prasad, N. 1998. Evolution of the GPGP Domain-Independent Coordination Framework. Computer Science Technical Report TR-98-05, University of Massachusetts at Amherst. A longer version of this paper will appear in the *Journal of Autonomous Agents and Multi-Agent Systems* in 2003.
- Lesser, V.; Horling, B.; and et al. The TÆMS whitepaper / evolving specification. <http://mas.cs.umass.edu/research/taems/white>.
- Raja, A.; Lesser, V.; and Wagner, T. 2000. Toward Robust Agent Control in Open Environments. In *Proceedings of the Fourth International Conference on Autonomous Agents (Agents2000)*.
- Shehory, O.; Sycara, K.; Sukthankar, G.; and Mukherjee, V. 1999. Agent aided aircraft maintenance. In *Proceedings of the Third International Conference on Autonomous Agents (Agents99)*, 306–312.
- Vincent, R.; Horling, B.; and Lesser, V. 2001. An agent infrastructure to evaluate multi-agent systems: The java agent framework and multi-agent system simulator. In Wagner, T., and Rana, O., eds., *Infrastructure for Agents, Multi-Agent Systems, and Scalable Multi-Agent Systems, Lecture Notes in AI*. Springer. 102–127.
- Wagner, T., and Lesser, V. 2001. Design-to-Criteria Scheduling: Real-Time Agent Control. In Wagner/Rana., ed., *Infrastructure for Agents, Multi-Agent Systems, and Scalable Multi-Agent Systems*, LNCS. Springer-Verlag. 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.
- Wagner, T.; Phelps, J.; Guralnik, V.; and VanRiper, R. 2004. Coordinators - coordination managers for first responders. In *To appear in the proceedings of the 3rd International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS04)*.
- Wagner, T.; Garvey, A.; and Lesser, V. 1998. Criteria-Directed Heuristic Task Scheduling. *International Journal of Approximate Reasoning, Special Issue on Scheduling* 19(1-2):91–118. A version also available as UMASS CS TR-97-59.
- Wagner, T.; Guralnik, V.; and Phelps, J. 2003a. 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. Nominated for most novel application award.
- Wagner, T.; Guralnik, V.; and Phelps, J. 2003b. 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.