
IDIS Laboratory Technical Report No. 105
**MADGS: An Architecture for Dynamic, Multi-Commander,
Multi-Mission Planning and Execution**

EUGENE SANTOS JR.^{a*}

SCOTT DELOACH^{b*}

MICHAEL T. COX^{c*}

SEPTEMBER 2003

Abstract

The Multi-Agent Distributed Goal Satisfaction (MADGS) project facilitates distributed mission planning and execution in complex dynamic environments with a focus on distributed goal planning and satisfaction. By understanding the fundamental technical challenges faced by our commanders on and off the battlefield, we can help ease the burden of decision-making. MADGS lays the foundations for retrieving, analyzing, synthesizing, and disseminating information to our commanders. In this paper, we present an overview of the MADGS system and discuss its key components that has formed our initial prototype and testbed.

^a Department of Computer Science & Engineering, University of Connecticut, Unit-155, Storrs, CT 06269-3155, eugene@cse.uconn.edu

^b Department of Computing and Information Science, Kansas State University, Manhattan, KS 66506-2302, sdeloach@cis.ksu.edu

^c Department of Computer Science & Engineering, Wright State University, Dayton, OH 45435-0001

* This project was supported in part by Air Force Office of Scientific Research Grant No. F49620-99-1-0244.

Contents

1	Introduction and Motivation	4
2	The MADGS Environment	5
3	Carolina Agent Server	7
4	Deploying MADGS	9
4.1	Multiagent Systems Engineering	9
4.2	agentTool	13
5	DGS	19
5.1	Intelligent Resource Allocation	22
5.2	Modeling Intra-Organizational Logistics	23
6	PRODIGY	26
6.1	Prodigy/Agent	26
6.2	GTrans	26
6.2.1	Goal Transformations	27
6.2.2	GTrans User Interface	28
7	Putting It All Together	29
7.1	Evaluation	30
7.1.1	GTrans Evaluation	30
7.1.2	Re-planning Study	32
8	Conclusion	35

List of Figures

2.1	MADGS System. Component/Agent Interaction Flow.	6
3.1	Carolina Architecture.	8
4.1	MaSE Methodology	10
4.2	Role Model	11
4.3	Concurrent Task Diagram	12
4.4	Agent Class Diagram	12
4.5	Agent Conversation Diagrams	13
4.6	Deployment Diagram	14
4.7	agentTool	15
4.8	Analysis to Design Transformations	15
4.9	Agent Transformation Architecture	16
4.10	Basic AgentComponent State Diagram	17
4.11	Mobile AgentComponent State Diagram	18
4.12	Error Messages	19
4.13	Deadlock Detection	19
4.14	Conversation Code Generation	20
6.1	The GTrans interface.	28
7.1	Military Interdiction Scenario. Rivers have target bridges spanning them.	30
7.2	Resource allocation and negotiation.	31
7.3	Plan execution commences after initial plan generated.	31
7.4	Goal satisfaction as a function of cognitive model.	32
7.5	Goal satisfaction as a function of problem complexity.	32
7.6	Goal satisfaction as a function of expertise.	33
7.7	Replanning time for standard (without GTrans) and focused (with GTrans).	34
7.8	Replanning nodes explored for standard and focused.	34

1 Introduction and Motivation

The strength of the US military is founded upon the autonomy and independence given to its battlefield commanders who effectively and efficiently carry out their assigned roles and tasks. Our commanders draw upon their wealth of training, personal resourcefulness, and sheer ingenuity in order to accomplish their missions. This is especially critical when faced with the constantly changing and adverse conditions on the battlefield. Thus, the ability of our individual commanders has been and still remains the key to US military dominance in the world.

Examining the broader picture of military operations, its success rests not only in the individual commander but also in the ability of our commanders to cooperate and adapt in a timely manner to achieve their goals. The need to adapt again arises from the changing conditions on the battlefield be it changes in overall mission or changes induced from goal failures.

In particular, for the US Air Force, real-time air mission planning and execution occurs in a highly complex and dynamic environment where new constraints and conditions frequently arise at all levels of operation from theatre level planning to individual unit tasking and execution. Some new conditions are “discovered” as planning failures arise while completing local objectives (goals and sub-goals). Planning failures are caused by factors such as internal planning conflicts, irrevocable commitments from partial plan execution, and insufficient resources (number of planes, fuel, etc.). Other, conditions are imposed from outside the system during planning and execution. Such conditions include command modifications in the mission objectives (human intervention – friendly and/or enemy) and continuous changes in the external environment (weather changes, equipment failure, etc.). Moreover, military operations planning is a distributed activity. Planners at all levels and in all services cooperate (and compete) to create and execute operations.

With all this in mind, the burden faced by our commanders/planners in the modern information rich battlespace is overwhelming to say the least. To be fully aware of the battlespace situation is simply not feasible. However, the strength of our commanders lies in what they know combined with what they can easily access in order to get the job done.

Unfortunately, most of the research to date on mission planning and execution has been based around the unrealistic assumption of a single “point of planning;” that is, all planning occurs at a single location with unlimited access to global information. Such an approach does not permit scaling and is unrealistic in today’s information intensive environments. First of all, without the ability to scale to large problems, automated planning will not be able to fit into existing organizations, or to permit simultaneous planning with plan execution. A particularly important issue is that local dynamic conditions, as discussed in the preceding paragraph, can result in a rippling effect of plan failures all the way up to the highest levels. Secondly, unlimited global access is simply not feasible given the size, reliability, and dynamic reconfigurability of our information network as well as the necessity of providing information security in such a critical environment. For mission planning and execution to work effectively, we must be able to reliably provide commanders the information when they “need to know” and guarantee the information security if they “need to know.”

In the real-world dynamic and large-scale setting of air mission planning and execution, this can only be achieved in a distributed manner to provide sufficient levels of efficiency, robustness, and security. This has been the primary theme for our Multi-Agent Distributed Goal Satisfaction (MADGS) project.

In our approach, we set out to clearly identify the root causes of plan failures and when and where such failures can be best addressed by the human commanders. In particular, we posited that significant planning failures occur during execution when specific resource requirements for a given task cannot be satisfied because of the dynamic nature of our operational environment. Instead of

forcing a costly re-planning of the mission, we redefined the problem in terms of satisfying resource requirements during mission execution. This allows us to avoid planning failures if alternative resources can be located (in a local fashion) for the human commanders to carry out their tasks. In the case where resources cannot be allocated, we now have up to date information on relevant resources and availability that can now be passed up the mission plan hierarchy. Combined with our developments in mixed-initiative planning through goal-transformations, we can better assist during re-planning but also limit the scope necessary of the overall re-planning effort. Specifically, we have achieved the following: We (1) developed and deployed a scalable mobile multi-agent infrastructure for dynamically reconfigurable networks; (2) designed a model of distributed goal satisfaction to mitigate plan resource failures; (3) built a framework for automated software agent generation and validation; and, (4) introduced mixed-initiative planning through goal transformations.

2 The MADGS Environment

The Multi-Agent Distributed Goal Satisfaction (MADGS) environment is a JAVA-based mobile-agent system that facilitates distributed mission planning and execution in complex dynamic environments with a focus on distributed goal satisfaction [39]. The MADGS system represents the union of five separate components, Agent-Server (named Carolina), mobile-agents, Distributed Goal Satisfaction (DGS), agentTool, and PRODIGY. The primary issues we explored include robust and reliable communication protocols, agent design, and a system architecture that facilitates both agent and agent server autonomy.

The target real-world operational environment for the MADGS system is a network topology that consist of intermittent nodes and uncertain network connections that exist in a large-scale, multi-platform dynamic network. The resulting design developed for this environment addresses the communications issues faced when handling massive numbers of mobile-agents in such a topology. Our development process required the consideration of bandwidth capacities (minimal broadcasts if any), mobile-agent collaboration issues, and server awareness of available resources. In designing a system capable of handling an unknown but unrestricted number of communication and agent migrations over the proposed topology we made an in- depth examination of both agent and server responsibilities. From this examination we developed a premise that there exists a marriage between the functionality of the Carolina agent-server and the agents themselves despite their autonomy. For this reason the MADGS architecture was built around this marriage, maintaining autonomy for both without depreciating the security or performance of the system. The marriage is one built of necessity. In order to minimize agent size some functionality was better placed in the server and offered as a service to the agents.

MADGS mobile-agents are the workhorses of the system providing the functionality the system users require. Mobile-agents are injected into the system through the agentTool [17] component responsible for agent creation. The agentTool component is not a standard component present on all nodes. As an administrative tool, agentTool, is instantiated at predetermined points in the network expressly as a defined creation and entry point for mobile-agents into the MADGS system. This architecture will provide an avenue for insuring a level of authenticity and thereby security to the system. If a new mobile-agent or a clone is needed, the requests for these are filtered through an agentTool component via an agentToolHandler agent. The DGS module will provide alternative resource configurations to facilitate the completion of a plan constructed via PRODIGY given constraint failure without backtracking or replanning. The DGS component will interact with the resource agent to maintain an accurate record of available resources using a distributed database of resource attributes and linear programming tools to make replacement determinations. The

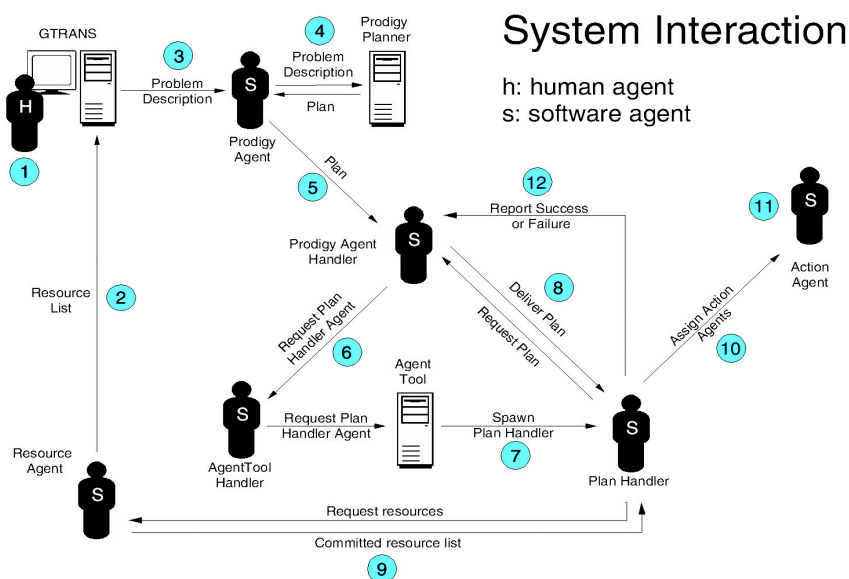


FIG. 2.1. MADGS System. Component/Agent Interaction Flow.

DGS component will rely on human operators to make resource substitution decisions. Alternative resource configurations for resources identified by the plan as mission-critical are given priority by the DGS component. PRODIGY [48] is a legacy planning system that will provide the initial master plan for operations. Figure 2.1 illustrates the interactions between agentTool, PRODIGY, Carolina, and the DGS components of MADGS. In this figure, Carolina is the backplane responsible for the communications and agent operation environment.

Our MADGS system functions as follows (Figure 2.1): For each human commander, there is an associated PRODIGY system complete with GTrans interface to assist the commander in formulating plans for their given missions. Also attached to the commander are a ResourceAgent, a PRODIGY agentHandler, and agentToolHandler. When a plan is completed by the commander with PRODIGY, it consists of a sequence of specific execution tasks. The PRODIGY agentHandler now requests agentToolHandler to spawn a specific planHandler agent from agentTool for this given plan. In this case, the planHandler can be retrieved from a library of agents that agentTool has created offline. The plan sequence is now picked up by our newly created planHandler agent where a resource analysis is conducted to determine the resource requirements of each task in the sequence. Execution of each task is now commenced. For each task being deployed, MADGS attempts to satisfy the specific resource requirement by querying the commander's ResourceAgent and then negotiating with other ResourceAgents "nearby" if necessary. When the requirements are satisfiable, the planHandler then requests agentToolHandler to now spawn a specific task execution agent for the task to be executed. This continues until all tasks are executed or a failure in a task execution or resource requirement is encountered. Once encountered, the planHandler then gathers the details of the failures and attempts to overcome the failure to send to the prodigyAgentHandler and mixed-initiative replanning commences. Also, implicit in the diagram is the ability of the human decision maker to observe activities on the system.

In the following sections, we detail the individual components used in MADGS in our above description.

3 Carolina Agent Server

For MADGS, the critical element for agent communications lies in the necessity of not having a *singular point of failure* in the system. This obviously means that we cannot afford to use centralized directories or look-up tables. Due to the network constraints we also cannot leave communications up to individual agents since this will effect the size of mobile agents and thereby bandwidth consumption. Another communications issue facing development of the MADGS system was location of agents and resources. Without centralized directories, look-up tables or the ability to farm communication concerns out to agent resources our research led to the quick determination that no available system answered the communication issues, resource tools or adaptability our project called for. For this reason we embarked on the design and implementation of the Carolina agent server and mobile agent system. Unfortunately no generic agent system with sufficient basic communications and resource management abilities could be effectively utilized.

Systems identified¹ made different assumptions and focused on a specific problem without offering some base functionality that all agent systems could build off of. It is hoped that one by-product of our work is to define the base functionality needed by all agent systems. This is not to say that all agent systems should or will be based off of our work given all systems do not work off general models.

The communication protocols we devised in Carolina attempts to insure several things:

1. The size of an agent will not increase significantly with increased interactions;
2. Each node will have a consistent view of the 'world';
3. The network load posed by communications can be minimized compared with alternate communication methods;
4. There is no central point of failure in the system; and,
5. All communications can be routed within a reasonable time frame.

These assurances can be made despite the volatile and intermittent nature of our network environment. The first assurance follows from the need for agents to only have an agent's unique assigned name to maintain a communication link with that agent. The second assurance is possible because our system provides for information sharing between nodes that guarantees all nodes will know the exact location of all mobile resources (agents are resources) assuming no non-server resource movement for a time-period not to exceed some constant value τ . The third assurance is possible because our communication protocol negates the need for broadcast except for system-wide alerts that are hypothesized to be rare in any system. Since our protocol does not utilize a central look-up table or centrally located directories there can be no single point of failure for communications thus the forth assurance. The last assurance stems from the second assurance, since all nodes are guaranteed to know the location of all mobile-agents it is therefore possible to route ALL communications within a reasonable time frame.

The foundation of the MADGS system is the agent server named Carolina. Carolina has a three-tier architecture with several internal components that are described in this section. There are four main functions of Carolina:

1. Provide an agent execution environment;
2. Insure system integrity through role-based security techniques;
3. Allow access to system resources where appropriate; and,
4. Provide communication services that improve the overall system performance.

¹We conducted a survey of all available agent systems and development environments. See <http://excalibur.brc.uconn.edu/madgs/agentsurvey.doc>.

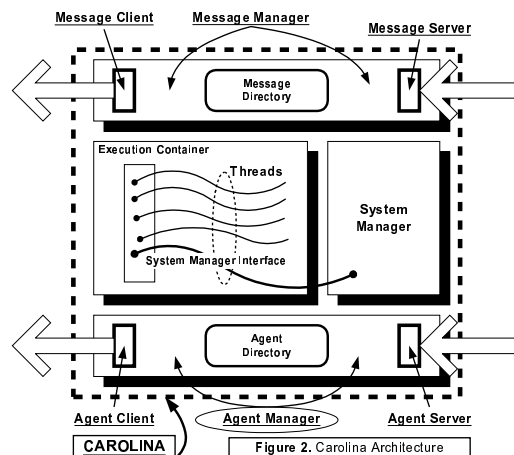


FIG. 3.1. Carolina Architecture.

The first function is basic to all agent systems and needs no further mention. Carolina System Manager Component oversees the agent interactions to insure only appropriate communications occur between different types of agents using case-based reasoning techniques. The second function controls the permissible access of resources by agents through role based security techniques.

The access Carolina provides to system resources to requesting agents. These resources may only be accessed through proper channels, in other words through the proper agent since agents represent all resources. For example, an attempt to access a database directly instead of through the databases interface agent would meet with failure and cause retribution from Carolina as a hostile act. In this event, the System Manager will sever the agents thread and report the offending agents class and offense to a system user. The last function of Carolina currently revolves around communications. There are three things of interest in this area, server-to-server communications, agent migration, and agent-to-agent communications. Server-to-Server communications that is crucial to the operation of the MADGS system however only occurs twice, upon Server startup and shutdown. The primary function of the server-to-server communication is the location and interconnection the MADGS system, Carolina servers and the agents they host. Agent migrations in Carolina are handled as a form of communication though migrations are handled through a separate port than communication messages. Since agents in the MADGS system are not allowed to directly communicate with another agent, Carolina must offer a service for local and remote communications.

The decision to not allow agents to directly communicate departs strikingly from the standard agent-based system protocols. The reason we take this unique stance is the limitations that occur when you allow an unlimited number of agents to freely migrate and communicate over a constantly changing network topology. Under such conditions the network traffic increases significantly as mobile-agents attempt to maintain the current location of other mobile-agents they are collaborating with. Given our network environment, bandwidth usage must be kept to a minimum; therefore management of this cannot be left to the individual agents. It becomes an issue of control and performance. Let us begin to examine the Carolina architecture (See Figure 3.1).

Carolina receives agents through it AgentServer Port. The AgentManager controls this port. The AgentManager receives incoming agents, checks their intended destination (IP), if it is local

then the AgentManager registers the agent in the AgentDirectory, deserializes it and passes it to the ExecutionContainer where the agent is provided with a thread for execution. If however, the agent's intended IP is not local, the AgentManager simply reroutes the serialized agent through the AgentClient Port after registering the transient agent in the AgentDirectory. The AgentDirectory is one of the key components of the Carolina communication scheme. It maintains data on all agents that the resident Carolina server has seen. Information stored in the AgentDirectory includes typical information including the agent's unique name assigned at creation, the agent's class, source IP, current (or last known) IP, and goal. Additionally, AgentDirectory stores a pointer to the messages stored in the Message Directory for individual agents.

Finally, since the MADGS system has been designed for use in a large-scale dynamic network architecture with massive numbers of mobile agents carrying out communications, point-to-point communications were not an option. The protocol we developed extends existing work [28]. Agents are not responsible for maintaining an address book of any kind. If communication is needed with a specific agent the agent only needs that agent's unique identification tag. Location of and routing of messages to agents in the network is performed as a joint effort between the servers and a Communications agent. The server logs all agents entering or passing through the server in route to another node in the network. Carolina logs the agent name, unique identification tag, class, location or destination and time-stamp. In conjunction with this logging activity Communications agents use a random algorithm to canvas the network moving from server to server collecting the server's agent directory and compare it to their 'view' of the world. This agent then modifies (or cleans) the servers agent directory making additions, deletions and correction as needed.

4 Deploying MADGS

To support MADGS, we created the *Multi-agent Systems Engineering* (MaSE) methodology and an agent development environment called *agentTool* [15, 3, 17]. Using agentTool, developers follow MaSE, which guides them, step by step, through the analysis and design of complex, distributed, and dynamic multiagent systems, such as distributed mission planning and execution systems. agentTool is a graphically based, interactive software engineering tool that fully supports MaSE. agentTool helps developers in specifying multi-agent organizations and then semi-automatically generating designs and correct, executable code. MaSE and agentTool are both independent of any particular agent architecture, programming language, or communication framework. Using agentTool, it is possible to generate implementations targeted at various frameworks without changing the design. To support the MADGS architecture, we have developed specific code generation modules to produce systems that work in the Carolina framework.

4.1 Multiagent Systems Engineering

The general flow of MaSE follows the seven steps shown in Figure 4.1. The rounded rectangles on the left side denote the models used in each step. The goal of MaSE is to guide a system developer from an initial system specification to a multiagent system implementation. This is accomplished by directing the developer through this set of inter-related system models. Although the majority of the MaSE models are graphical, the underlying semantics clearly and unambiguously defines specific relationships between the various model components.

MaSE is designed to be applied iteratively. Under normal circumstances, we would expect the developer to move through each step multiple times, moving back and forth between models to ensure each model is complete and consistent. While this is common practice with most design

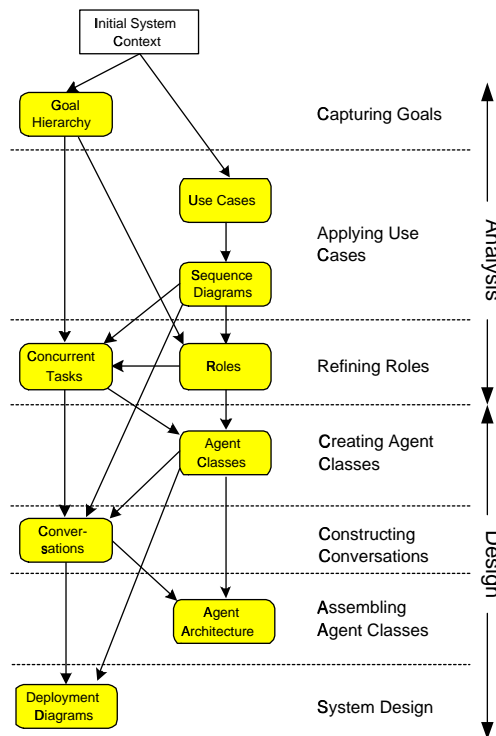


FIG. 4.1. MaSE Methodology

methodologies, MaSE was specifically designed to support this process by formally capturing the relationships between the models. By automating the MaSE models in agentTool, these relationships are captured and enforced thus supporting the developer's ability to freely move between steps. The result is consistency between the various MaSE models and a system design that satisfies the original system goals. The two phases of MaSE, Analysis and Design, are discussed in more detail below.

MaSE Analysis The Analysis phase includes three steps: capturing goals, applying use cases, and refining roles. In the Design phase, we transform the analysis models into models useful for actually implementing the multiagent system. Each of these steps are described below.

The first step in the MaSE methodology is Capturing Goals, which takes the initial system specification and transforms it into a structured set of system goals. There are two steps to Capturing Goals: identifying the goals and structuring goals. Goals are identified by defining the main purposes of the system. Once the system goals have been captured and explicitly stated, they are less likely to change than the detailed steps and activities involved in accomplishing them [24]. After identification, the goals are analyzed and structured into a Goal Hierarchy Diagram, which is an acyclic directed graph where child nodes are subgoals of the parent goal. There is typically a single system goal, which is decomposed into a set of subgoals. These subgoals are assigned to roles, and eventually to agents. Thus agents, based on the role they are playing, become responsible for achieving specific system goals.

The Applying Uses Cases step is crucial in translating goals into roles and associated tasks. Use cases are drawn from the system requirements and describe sequences of events that define desired system behavior. Use cases are examples of how the system should behave in a given case. To help determine the actual communications in a multiagent system, the use cases are converted into Sequence Diagrams. Sequence Diagrams depict sequences of events between multiple roles and, as a result, define the communications that must exist between the agents playing those roles in the final design. The roles identified here form the initial set of roles used in the next step. The events are also used later to help define tasks and conversations.

The third step in the Analysis phase is to identify all roles, starting with the set defined in the previous step, and to define role behavior and communication patterns. A role is an abstract description of an entity's expected function and is similar to the notion of an actor in a play [24]. Roles are identified from the Sequence Diagrams developed during the Applying Use Cases step as well as the system goals defined in Capturing Goals. MaSE ensures that all system goals are accounted for by associating each goal with a specific role, which is eventually played by at least one agent in the final design. Roles are captured via Role Models as shown in Figure 4.2. The boxes denote roles, which include a list of goals assigned to that specific role. Each role has at least one task, which are denoted by attached ovals. Communications between tasks are denoted by arrows pointing from the initiating task to the responding task.




FIG. 4.2. Role Model

Once roles have been identified, concurrent tasks are created to define the expected role behavior. Concurrent tasks are captured via finite state models and specify a single thread of control that integrates inter-agent as well as intra-agent communications. Task execution is based on its type. If a task is persistent, it starts when the agent is created and runs until the agent dies or the task reaches an end state. Persistent tasks are identified by having null transitions from the start state to another state in the concurrent task diagram. If a task is transient, the task is started in reaction to an incoming event. Transient tasks are recognized by having an incoming event on the transition from the start state. An example of a MaSE Concurrent Task Diagram is shown in Figure 4.3.

Role behavior is captured by a set of n concurrently executing tasks (where $n > 0$). Activities are used to specify actual functions carried out by the agent and are performed inside the task states. While these tasks execute concurrently and carry out high-level behavior, they can be coordinated using internal events. Internal events are passed from one task to another and are specified on the transitions between states. To communicate with other agents, external messages can be sent and received. These are specified as internal send and receive events. The syntax associated with state transitions is as follows

$$\text{trigger}(args1) [\text{guard}] / \text{transmission}(args2)$$

This is interpreted to mean that if an event *trigger* is received with a number of arguments *args1* and the condition *guard* holds, then the message *transmission* is sent with the set of arguments *args2*. Actions may be performed in a state and are written as functions. Besides communicating with other agents, tasks can interact with the environment via reading percepts or performing operations

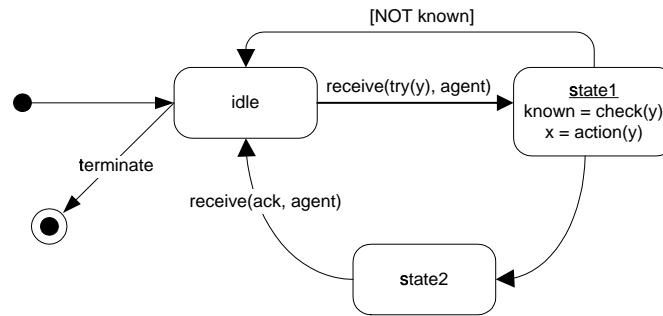


FIG. 4.3. Concurrent Task Diagram

that affect the environment. This interaction is typically captured by functions executed while in specific states.

MaSE Design Once the concurrent tasks of each role are completely defined, the MaSE Analysis Phase is completed and design begins. The MaSE Design Phase is used to create agent classes (from which agents are instantiated), the conversations between agent classes, and the internal definition of the agent classes. This is accomplished in four steps: Creating Agent Classes, Constructing Conversations, Assembling Agents, and System Deployment.

In the Creating Agent Classes step, agent classes are created based on roles from the analysis phase. Agent classes and the conversations between them are captured via Agent Class Diagrams, Figure 4.4, which depicts agent classes as boxes and the conversations between them as lines connecting the agent classes. Each agent class is assigned to play at least one role while multiple agent classes may be assigned the same role. Since agents inherit the communication of their roles, any communications between roles become conversations between their respective classes. Thus, as roles are assigned to agent classes, the overall organization of the system is defined.




FIG. 4.4. Agent Class Diagram

Once the conversations have been identified, the next step, Constructing Conversations, which define coordination protocols between two agents. Specifically, a conversation consists of two Communication Class Diagrams, one each for the initiator and responder. A Communication Class Diagram is a pair of finite state machines that define a conversation between two participant agent

classes. Both sides of a conversation is shown in Figure 4.5. The initiator always begins the conversation by sending the first message. The state machines should be consistent with each other; all messages sent by one side of the conversation should be able to be received by the other with reaching deadlock. The syntax for Communication Class Diagrams is similar to that of Concurrent Task Diagrams. The main difference between conversations and concurrent tasks is that concurrent tasks may include multiple conversations between many different roles and tasks whereas conversations are binary exchanges between individual agents.

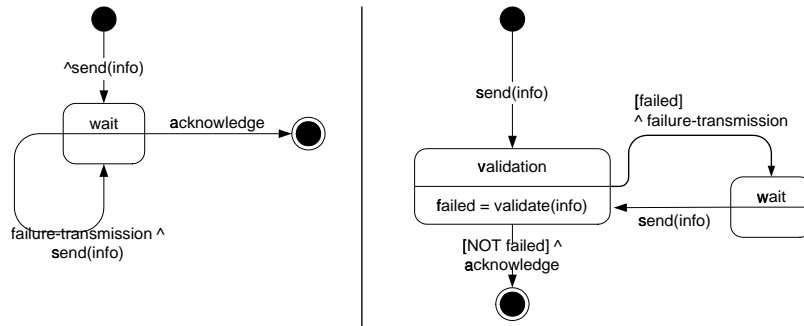


FIG. 4.5. Agent Conversation Diagrams

In the Assembling Agents step, the internals of agent classes are created. Robinson [38] describes the details of assembling agents from a set of standard or user-defined architectures. This process is simplified by using an architectural modeling language that combines the abstract nature of traditional architectural description languages with the Object Constraint Language, which allows specification of low-level details. The actions specified in the tasks and conversations must be mapped to internal functions of the agent architecture, while a link between actions and conversations must also be made.

The final step of MaSE is the System Deployment, in which the configuration of the actual system to be implemented is defined. In MaSE, the overall system architecture is defined using Deployment Diagrams to show the numbers, types, and locations of agents within a system as shown in Figure 4.6. The three dimensional boxes denote individual agents while the lines connecting them represent actual conversations. A dashed-line box encompasses agents that are located on the same physical platform.

The agents in a Deployment Diagram are actual instances of agent classes from the Agent Class Diagram. Since the lines between agents indicate communications paths, they are derived from the conversations defined in the Agent Class Diagram as well. However, just because an agent type or conversation is defined in the Agent Class Diagram, it does not necessarily have to appear in a Deployment Diagram.

4.2 agentTool

The agentTool system is our attempt to implement a tool to support and enforce MaSE. Currently agentTool implements all seven steps of MaSE as well as automated support for transforming analysis models into design models. The agentTool user interface is shown in Figure 4.7. The menus across the top allow access to several system functions, including a persistent knowledge base [37] conversation verification [31] and code generation. The buttons on the left add specific




FIG. 4.6. Deployment Diagram

items to the diagrams while the text window below them displays system messages. The different MaSE diagrams are accessed via the tabbed panels across the top of the main window. When a MaSE diagram is selected, the designer can manipulate it graphically in the window. Each panel has different types of objects and text that can be placed on them. Selecting an object in the window enables other related diagrams to become accessible.

We approached automated agent synthesis along three major themes in agentTool: (1) automated transformation of analysis models into design models, (2) verification of analysis and design models prior to code generation, and (3) code generation for a variety of agent frameworks.

Analysis to Design Transformations This section describes how agentTool transforms the analysis models into design models. Figure 4.8 shows the analysis-to-design transformation process in agentTool. The goal of these transformations is to take the analysis specification and transform it into a consistent design. The process is semi-automatic in that the designer initiates the transformation process and guides it along when help is required. The initial input to the transformations is the set of agent classes and the roles assigned to each agent class. The transformation system then generates the agent conversations as well as the internal agent component-based design based on the role model and tasks from the analysis phase.

In the semi-automatic transformation system, the architecture shown in Figure 4.9 is used. The AgentComponent is the overall controller of the agent and is responsible for instantiating other components, routing external messages, and handling agent mobility. Each concurrent task is transformed into a sub-component of the AgentComponent. These sub-components have their own state diagrams and conversations.

The MaSE methodology makes it clear that an agent class' roles, in conjunction with the protocols between the tasks, determine each agent's conversations. Therefore, the transformations create a separate component for every task in each role that an agent is assigned to play. The concurrent task definition is then copied into the associated component state diagram. Next, the states and transitions belonging to conversations are extracted and replaced with actions that represent the execution of the conversation. Using this approach, the component's state diagram retains the coordination and internal events necessary to ensure the behavior of the component matches the

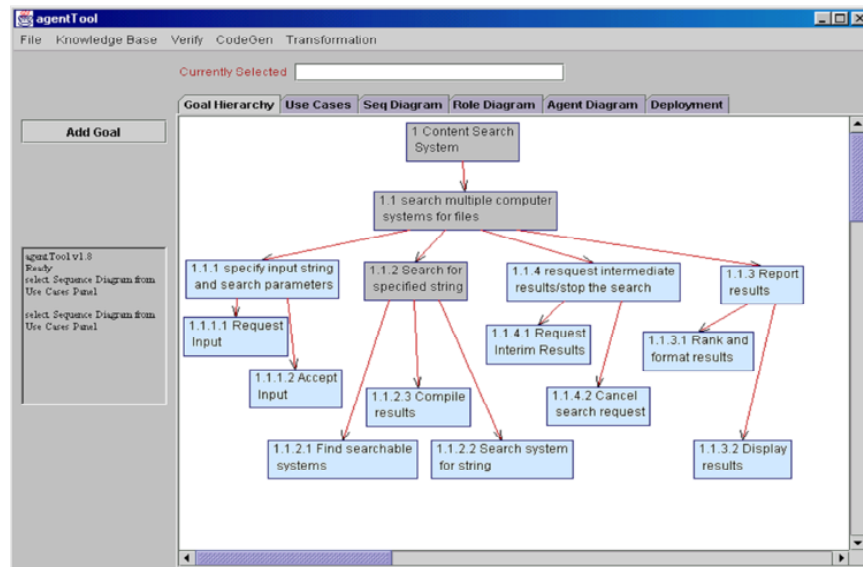


FIG. 4.7. agentTool




FIG. 4.8. Analysis to Design Transformations

task from which it was derived. Ultimately, the roles assigned to an agent determine the agent's components as well as the conversations in which it participates. Besides components derived from concurrent tasks, the transformations also create a Agent Component for each agent that captures intra-agent coordination. Figure 4.10 shows the state diagram for an Agent Component. The AgentComponent controls the instantiation of components and is responsible for handling conversation initiation messages received from other agents. The AgentComponent must determine if the message belongs to an existing component or if the AgentComponent must create a new component to handle the message.

The analysis-to-design transform is actually a sequence of transformations that incrementally change roles and tasks into agent classes, components, and conversations. Before beginning the analysis-to-design transformation process, the Role Model and its set of concurrent tasks, and the assignment of roles to agent classes must exist. During the first stage of the transformation process, agentTool derives agent components from their assigned roles and assigns external events to specific protocols. In the second stage, agentTool annotates the component state diagrams to determine where conversations start and end. During the last stage, agentTool extracts the

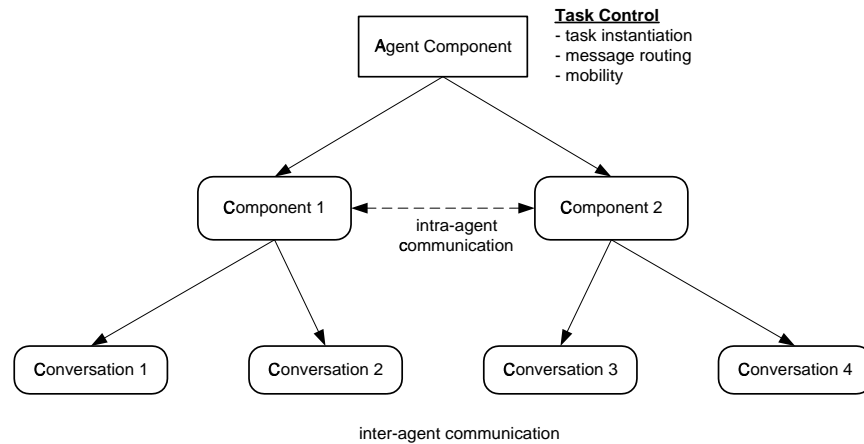


FIG. 4.9. Agent Transformation Architecture

annotated states and transitions and uses them to create new conversations, replacing them in the component state diagram with actions initiating the conversation. Detailed information, including a formal specification of the transformations, can be found in [46].

A second set of transformations currently implemented in agentTool consists of transformations to add functionality required for mobility. In the analysis phase, mobility is specified using a *move* activity in the state of a concurrent task diagram. This *move* activity is copied directly into the associated component state diagram during the initial set of analysis-to-design transformation described above. During the mobility transformation, the existing design must be modified to coordinate the mobility requirements between all components in the agent design. According to the mobility design approach, the AgentComponent is responsible for coordinating the entire move and working with the external agent platform to save its current state and actually carry out the move.

Specifically, the AgentComponent takes *move* requests from components and determines if the agent is ready to move. If the AgentComponent accepts a *move* request, it informs all other components, requesting their current state information. Once the other components send their state information and terminate, the AgentComponent requests a *move* from the agent platform. Thus, the AgentComponent is the repository for all component state information. After the agent has moved, the AgentComponent has the responsibility of restarting its components at the new location with their saved state information. Figure 4.11 shows the AgentComponent from Figure 4.10 transformed to handle the mobility requirements described above (the original parts of the state diagram are in the shaded area). This transformation is entirely automatic.

The second part of the mobility transform modifies the rest of the components in mobile agents. For a component that actually requests a *move*, the transformation replaces the original *move* activity with an internal *move* request event that is sent to the AgentComponent. The transform also adds the functionality to save the component's current state after it requests the move. All components in a mobile agent (not just those that request moves) must be able to respond to a *move required* message from the AgentComponent, save its state, and restart itself in the right state after the move. Because this process can occur in any or all states in the component state table,

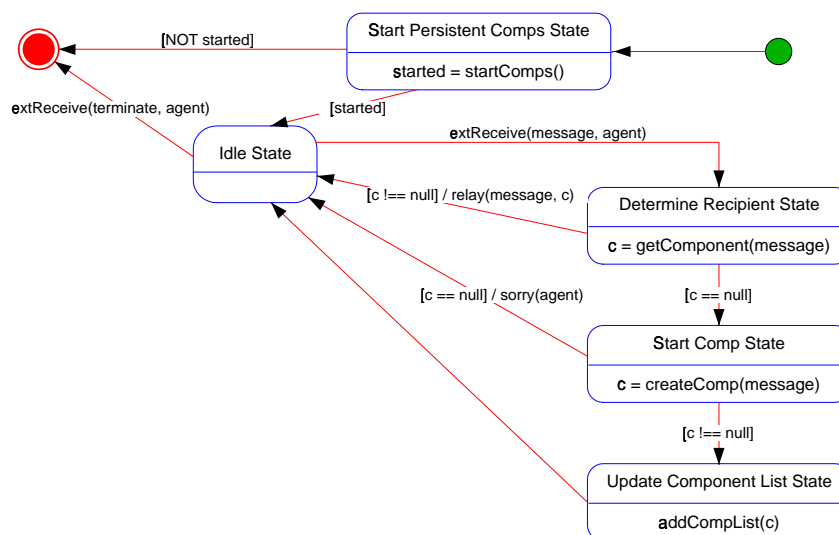


FIG. 4.10. Basic AgentComponent State Diagram

agentTool asks the designer to choose the states where a move would be allowable. More detailed explanations of the design-to-design mobile agent transformations are contained in [43].

Automated Verification The second research area in multiagent system synthesis was the automatic verification of agent protocols. Since it is critical that distributed mission planning systems such as those envisioned by MADGS operate properly without extensive testing, efforts were focused on ensuring that the protocols used for agent communications were as reliable as possible before implementation. agentTool, via the Spin system, currently checks classic conversation centric errors including deadlocks, non-progress loops, livelock, infinite overtaking, unused messages, and mislabeled transitions. agentTool also has the capability to verify that Sequence Diagrams generated during the analysis phase can actually be generated by the system design [31, 30].

agentTool performs these verification tasks by generating Promela code from the system design and passing that code to Spin. Spin then creates an analyzer to search the conversation state space, simulating all possible combination of messages in the conversation until either a deadlock condition occurs or the state space is exhausted. Conversations are considered deadlocked if they terminate in any state other than the end state. If a deadlock condition is detected, the analyzer writes a trace file that can be used to create a message sequence trace pinpointing the series of message events that led to the deadlock.

To detect non-progress loops, agentTool marks all states in the Promela conversation definition with the keyword *progress*, which is used by Spin to check that all states are entered during at least one execution path. If a state is not entered into then a non-progress error is generated. There are many causes for non-progress errors including livelock, infinite overtaking, deadlock, unused states and unused transitions.

Finally agentTool tests for valid message sequences by defining the message sequence in a *never* claim and checking to see if the sequence exists. If the sequence exists, Spin generates a never claim violation error. However, this is not really an error since the sequence is supposed to exist. If Spin does not report a never claim violation, the message sequence could not be found and even though


```
pan: invalid endstate (at depth 5)
pan: wrote verify.trail
(Spin Version 3.2.4 -- 10 January 1999)
Warning: Search not completed
  + Partial Order Reduction
Full statespace search for:
  never-claim          - (none specified)
  assertion violations +
  cycle checks         - (disabled by -DSAFETY)
  invalid endstates +
State-vector 24 byte, depth reached 8, errors: 1
  6 states, stored
  1 states, matched
  7 transitions (= stored+matched)
  1 atomic steps
hash conflicts: 0 (resolved)
(max size 2^18 states)
1.493    memory usage (Mbyte)
```

FIG. 4.12. Error Messages

```
DEADLOCK CONDITION EXISTS IN THE FOLLOWING CONVERSATION:
Conversation Name = SendInfo
Participant Name = Responder
Current State = validation
State Transition = null
DEADLOCK CONDITION EXISTS IN THE FOLLOWING CONVERSATION:
Conversation Name = SendInfo
Participant Name = Initiator
Current State = wait
State Transition = failureTransmission
```

FIG. 4.13. Deadlock Detection

performance of mobile and static multiagent systems in [35].

5 DGS

To recap, MADGS is an object-oriented system for the deployment of mobile-agents to facilitate large-scale planning and execution operations in domains such as manufacturing, search and rescue, and military operations. The general problem with the development of the MADGS system revolves around facilitation of real-time operation and plan failure handling. One of the common threads across the aforementioned domains and the general problem is the need for distributed goal satisfaction that can work cooperatively with legacy planning systems yet autonomously handle changes in constraints. The ability to autonomously handle changes in the constraints of a plan can mean the success or failure of any distributed operational mission/goal. The need to re-plan or backtrack due to constraint changes in any plan can mean a substantial resource loss; be it lost

```
Message m = new Message();
int state = 0;
boolean notDone = true;
/* state constant definitions */
final int StartState = 0;
final int State3 = 1;
final int State4 = 2;
final int StartState_out = 3;
final int State4_out = 4;
final int State3_out = 5;
.
.
.
while (notDone)
{ switch (state)
  { case StartState :
    state = StartState_out;
    break;
    case StartState_out :
    if (m.performative.equals("send"))
    { info = m.content;
      state = State4; }
    break;
    case State3 :
    state = State3_out;
    break;
    case State3_out :
    m = readMessage(input);
    if (m.performative.equals("send"))
    { info = m.content;
      state = State4; }
    break;
    case State4 :
    state = State4_out;
    break;
    case State4_out :
    if (! failed)
    { m = new Message();
      m.performative = "acknowledge";
      sendMessage(m, output);
      notDone = false; }
    if (failed)
    { m = new Message();
      m.performative = "failure-transmission";
      sendMessage(m, output);
      state = State3; }
    break;
  }
}
```

FIG. 4.14. Conversation Code Generation

capital or life, the expense is real. Our approach seeks to mitigate a significant amount of this loss by preemptively expecting failure, defining alternative constraint configurations, developing delivery arrangements and in the event of a failure offering an instant solution to the user.

This distributed goal satisfaction (DGS) process is a primary background activity of the MADGS system. In the forefront, the system is providing an environment for general operation, sub-plan and sub-task execution, and user interface. It is this aspect of the MADGS system that facilitates the DGS process. The MADGS system represents every resource (including personnel) by at least one agent. However, this representation allows the MADGS system to facilitate the DGS process by maintaining a more complete view of the current state of the 'world'. This world-view is constantly shifting in any operation especially large-scale operations. For this reason alone communications becomes a crucial aspect of our system.

Even though some agent-based mission planning and execution systems have been developed [49] they do not fully use the power of agent programming. Our approach is to use the strengths of legacy systems in conjunction with the strengths of agent programming coupled with our own approaches to communications (as outlined above) and resource location and allocation. Most large-scale operations create a plan off-line by formulating a problem or suggested outcome and then determining an optimal plan for the realization of this global goal. While an agent system could be used for such an operation there are legacy systems in existence optimized for this purpose. This is why PRODIGY is used to create a Master Plan for the MADGS system. This plan is then provided as input to the MADGS system. The MADGS system then uses command agents to decomposes the plan (if necessary) into sub-plans which are further decomposed into tasks by sub-command agents (Command agents with a lower rank) that assign the tasks to subordinate task agents. This process is not dissimilar to those present in existing agent-systems. Our approach is only unique in how we communicate and plan for alternative courses of action in the event a plan fails during execution due to changes in the present state of the 'world'.

Resolving to one course of action (or plan) in a real-time system poses great difficulty due to the volatile nature of the constraints and the conditions a plan is based on. A change in a constraint or condition of a sub-plan could lead to its total or partial failure that in turn can lead to a rippling effect, thereby negating the validity of the initial plan. To overcome these points of failure, a robust and flexible planning system is needed. The DGS agent module seeks to provide a surrounding technique to improve the robustness and flexibility of the overall planning and execution process.

We accomplish this by acting on the resources required to accomplish a given goal, plan or task. A resource is any commodity that is necessary to facilitate the completion of a goal. (i.e.: Goal A requires resource X, quantity 3) The DGS agent module receives the local version of the plan and a list of the required constraints (primary resources). With this input a tree is constructed of the alternative resource configurations meeting a stated Tolerance level representing a cost-benefit function for each required primary resource. This data is then rated based on alternative resource availability to the local plan (taking into consideration other known pending or current local plans). This information is then stored and the DGS agent module monitors the primary resource statistics. In the event that a primary resource fails or is exhausted the DGS agent module suggests alternative resource configurations to complete the current plan. If accepted the resources are set into action in place of the primary (failed) resources. During this process DGS collects and records data on the selections that users make to give weight to certain configurations and more importantly to learn new resource alternatives and configurations when users manually manipulate the suggested configurations prior to commitment. This process when successful negates the need to replan or backtracking furthering the maintenance of a real-time system. In the following subsection, we present our formal approach to resource matching.

5.1 Intelligent Resource Allocation

The need to re-plan or backtrack due to plan failure can mean a substantial resource loss. Be it lost capital or life, the expense is real. For MADGS, we determined that we can mitigate a significant amount of this loss by preemptively expecting failure, defining alternative resource constraint configurations, developing delivery arrangements.

In the event of a failure, our previous goal was to offer alternative solutions to the human commander in an effort to assist them in accomplishing their mission. Our approach sought to mimic and exploit the strengths of our on-site commanders by enhancing their own innate resourcefulness. In effect, we attempted to provide targeted resource information to assist commanders in resource substitution decisions and on the spot plan alterations.

At it's most basic level, resource substitution can be simplified to: Can resource A be substituted by resource B given a set of mission requirements. While this is the minimum that is needed to assist the commander's "scrounging" decisions, we realized that many more factors must be captured. These include cost factors such as transportation costs, scheduled availability, production costs, etc. Also, another more problematic cost is the possibility of cascading plan failures when a particular critical resource is diverted towards solving another plan. This is especially unacceptable if the second plan is not critical to overall theatre operations.

In addition, given the vast number of resources available in any theatre, the number of alternative resource suggestions can actually overwhelm the human commander. Our goal is to take all these issues into consideration and develop a framework for intelligent resource substitution that can ultimately further provide global guarantees of mission satisfiability in the overall theatre operations.

Clearly, resource substitution is very much related to logistics management (eg., [45]). We can state such a problem as consisting of a set of suppliers and consumers. Suppliers provide resources, while consumers utilize some resources to achieve some goals, such as doing jobs or producing final products. After more than a decade's development on logistics management, many kinds of logistics management models have been proposed and implemented. Some models are stand-alone and centralized, while others use a client/server approach([23, 4, 36]). In recent years, researchers have proposed multi-agent based models ([29, 51, 50, 32, 40]). However, most of the models regard logistics management as an auction, in which each entity tries to maximize its own benefit. Such an approach is only appropriate for inter-organizational logistics, which consists of competitive entities. However, this is clearly inappropriate for problem. In our case, suppliers (bases, depots, etc.) and consumers (commanders) have a common goal to maximize the outcome of the entire organization (theatre operations). People usually use intra-organizational logistics to describe this case. The point is that the maximized outcome doesn't mean the maximal benefit of each entity. Therefore, an optimal scheduling may be built based on sacrificing of some individuals' benefits. Unfortunately, the typical auction approach does not account for necessary self-sacrifice.

To address intra-organizational logistics, some researchers have developed coordinating multi-agent logistics management models. N.M. Sadeh et al. in [40] proposed MASCOT (Multi-Agent Supply Chain cOordination Tool), a reconfigurable, multilevel, agent-based architecture for coordinated supply chain planning and scheduling. W. Shen et al. in [44] proposed MetaMorph II architecture for enterprise integration and supply chain management, which is mediator-centric and agent-based. M.S. Fox et al. in [22] described the architecture of the integrated supply chain management system, in which each agent performs one or more supply chain management functions, and coordinates its decisions with other relevant agents. A KQML-based multi-agent coordination language was proposed in [1] for distributed and dynamic supply chain management. However,

their approaches are ad hoc and lacking in precise optimization models.

To overcome their limitations, we believe that it is necessary to formulate the resource allocation problem formally. Luh et al. in [33] utilized Lagrangian Relaxation to remove couplings between constraints so that the original problem can be separated into subproblems. Ideally, these subproblems should be separable/independent. The separability property is good for us because we can allocate a different agent to solve each subproblem. If the solutions for these subproblems are compatible with each other, we are done. Otherwise, these agents can exchange information and find an optimal way to satisfy constraints. We believe that the formulations we have developed before can satisfy separability since our formulations are similar to those found in [33] for manufacturing scheduling. It was demonstrated that the separability condition does hold for those problems. In case the separability does not hold in general, we can still significantly benefit from the problem decomposition and subproblem groupings to improve our computations.

5.2 Modeling Intra-Organizational Logistics

Before we can model intra-organizational logistics, we need to know its specific issues. Since suppliers (bases and other commanders) can only provide limited resources for any given period, the needs of the consumers (commanders) may not be fully satisfied. The goal of intra-organizational logistics is to reasonably allocate resources so that the profit of the whole organization is maximized. Depending on the following factors, the problem may be relatively easy or very complex:

- **Number of resource types:** For example, typical resource types are plane, ship, truck, fuel, ammunition, soldiers, airport, and so on.
- **Number of suppliers for each consumer:** If each consumer has only one supplier, the problem becomes easy. This is generally the case when a supplier is in charge of one geographical area. But with modern transportation means and networking, a supplier is no longer limited by its location. Therefore, the typical relationship between consumers and suppliers is a many-to-many relation.
- **Task properties:** A task can be the mission activities in a given period. Generally, a deadline is set on each task. The execution of a task needs a certain amount of resources. Sometimes, a task cannot be started unless all resources have been received. In other cases, lack of resources only delays the execution time or degrades the quality of a task. Similarly, extra resources may accelerate the execution of a task or achieve a better goal. A task may be viewed as a single step operation. Or it can consist of a sequence of stages. It is possible that some stages are critical. If a critical stage violates the deadline or cannot be continued for some reason, the whole task may be regarded as having failed. A consumer may only execute one task over a long time period or many tasks. Several tasks may be related, such as having a common goal. Related tasks have more constraints. For example, two tasks are required to begin at the same time or one before the other.
- **Resource properties:** Different resource types are not fully separated. Two different types of resources may have overlapping functionalities. For example, two kinds of planes can achieve the same objectives, only with different costs. We call this property resource exchangeability. Based on resource exchangeability, we can allocate alternative resources if critical resources are in demand. However, if alternative resources lead to higher costs, a trade-off exists between using these alternative resources and waiting for needed resources.
- **Uncertainty:** That the situation is changing over time increases the difficulty of managing intra-organizational logistics. For instance, tasks arrive dynamically because of great variability of customer demands and resources may not be provided in time. Due to the insufficient

resources or other uncertainties, like machine breakdown, tasks may not be finished before the deadline. There are so many uncertainties, how can we manage intra-organizational logistics efficiently?

Traditional models only consider subsets of these factors and deal with the problem in a centralized manner, which means all the information is collected at one place for analysis.

Our approach is different from existing approaches above in that we take into account a precise optimization model.² Through the use of Lagrangian Relaxation, we can decompose a resource allocation problem into subproblems, each of which will be solved by a specific agent. If the solutions for these subproblems are compatible with each other, we are done. Otherwise, these agents can exchange information with each other until a global optimal solution is found. Below, we provide various models for use depending on the complexity of the target planning problem.

Single Resource Type/Single Supplier/Single Job. In this case, we suppose that only one kind of resource exists in the system, and each consumer requests resources from a corresponding supplier to do a single job. We also fix the cost of using unit resource. Therefore, there is no need to consider the cost when we do job scheduling and the goal is to minimize job delay.

Since each consumer is related with only one job, we consider only jobs and suppliers in the remaining section. Let the number of jobs be N_c and the number of suppliers be N_s . For each job $i = 1, \dots, N_c$, b_i represents the starting time, c_i represents the completion time, p_i represents the execution time, d_i represents the deadline for job i , s_i represents the supplier from which job i will get resources ($s_i \in [1, N_s]$), and r_i represents the needed resource number. Also, we use T_i to measure the delay of job i ($T_i \in [0, c_i - d_i]$). Because some jobs may be more important than others, it is not desirable to delay these jobs. To reflect this factor, we use w_i to represent the importance of each job. The higher the value of w_i is, the more important the job is. We can define the objective function as:

$$\text{minimize } \sum_{i=1}^{N_c} w_i T_i^2$$

At any given time t , the total granted resources from one supplier must be less than its capacity. We use N_j ($j \in [1, N_s]$) to represent the amount of resources that supplier j has. We suppose T is long enough to complete all the jobs. The resource capacity constraints can be expressed as:

$$\sum_{i=1}^{N_c} \delta_{ijt} \cdot r_i \leq N_j, j = 1, \dots, N_s, t = 1, \dots, T$$

$$\delta_{ijt} = \begin{cases} 1 & \text{If } s_i = j \wedge b_i \leq t \leq c_i \\ 0 & \text{Otherwise} \end{cases}$$

In addition, the following processing time constraints must be satisfied:

$$c_i - b_i + 1 = p_i, i = 1, \dots, N_c$$

Single Resource Type/Single Supplier/Multiple Jobs. In this case, each consumer can execute several jobs. We assume that at any time, each consumer can execute at most one job. Therefore, jobs belonging to one consumer cannot be overlapped. Here, we suppose that s_{ik} represents minimal switching time for executing job k after finishing job i if these two jobs are executed

²Existing approaches include [22, 53, 40, 52, 32, 22, 50].

by the same consumer. For each job i , o_i represents the consumer who will do the job. Compared to the above case, the precedence constraints need to be considered:

$$\delta_{ik}(c_i + s_{ik} + 1 - b_k) + (1 - \delta_{ik})(c_k + s_{ki} + 1 - b_i) \leq 0,$$

where $i, k = 1, \dots, N_c, i \neq k, o_i = o_k$;

$$\delta_{ik} = \begin{cases} 1 & \text{If job } k \text{ occurs after job } i \text{ has been finished} \\ 0 & \text{Otherwise} \end{cases}$$

Single Resource Type/Multiple Suppliers/Multiple Jobs. The drawback of the above models is that a consumer cannot execute a job until the corresponding supplier has enough resources. Through allowing a consumer to request resources from multiple suppliers, this model is more flexible. In most cases, this model can get better job scheduling than the above two models. However, this model is more complex. Though we can limit each consumer to only requesting resources from a specific set of suppliers, we assume each consumer can request resources from all suppliers. But the costs of using resources from different suppliers are different. Therefore, we need to consider the cost factor in the objective function, shown as the following:

$$\text{minimize } \sum_{i=1}^{N_c} (w_i T_i^2 + \sum_t \sum_{j=1}^{N_s} r_{ijt} c_{rj}),$$

where r_{ijt} represents the number of resources that job i gets from supplier j at time t ; c_{rj} ($j \in [1, N_s]$) is the cost of using unit resource from supplier j per unit time.

Also we modify the resource capacity constraints:

$$\sum_{i=1}^{N_c} r_{ijt} \leq N_j, j = 1, \dots, N_s, t = 1, \dots, T$$

$$\sum_{j=1}^{N_s} r_{ijt} = \begin{cases} r_i & \text{If } b_i \leq t \leq c_i \\ 0 & \text{Otherwise} \end{cases}$$

Multiple Resource Types/Multiple Suppliers/Multiple Jobs. If we permit multiple jobs to be executed by one consumer, we need to consider the precedence constraints at the same time:

$$\delta_{ik}(c_i + s_{ik} + 1 - b_k) + (1 - \delta_{ik})(c_k + s_{ki} + 1 - b_i) \leq 0,$$

where $i, k = 1, \dots, N_c, i \neq k, o_i = o_k$;

$$\delta_{ik} = \begin{cases} 1 & \text{If job } k \text{ occurs after job } i \text{ has been finished} \\ 0 & \text{Otherwise} \end{cases}$$

The objective function and resource capacity constraints are the same as those in the fifth model.

In summary, we have deployed this approach using the various models above in a testbed currently separate from MADGS. We have achieved very good efficiency results from our testing and simulation. Our technique provides the ability to guarantee that the resource decisions made by the commander while satisfying the commander's need will also maximize the overall operational

success in the theatre without the need for centralized logistics optimization. In effect, we can naturally decompose and distribute the resource decisions for effective computation. This also allows us to provide the commander with informed knowledge concerning the expected impacts of such resource substitution decisions.

Our next step will be to deploy this intelligent resource allocation strategy within MADGS. For complete details of this model and experimental results, see Santos et al. [41, 42].

6 PRODIGY

For this project, we developed a number of computerized systems (1) to test our theory of distributed goal satisfaction (goal transformation theory), (2) to provide an interface for mixed-initiative team planning (distributed planners; some human, some machine), and (3) to wrap the main autonomous planner, the PRODIGY planning architecture, with a veil that behaves as an autonomous agent([2, 6, 8, 11, 18, 26, 27, 12, 47]).

6.1 Prodigy/Agent

PRODIGY [5, 48] is a legacy system that employs a state-space nonlinear planner and follows a means-ends analysis backward-chaining search procedure that reasons about both multiple goals and multiple alternative operators from its domain theory appropriate for achieving such goals. A domain theory is composed of a hierarchy of object classes and a suite of operators and inference rules that change the state of the objects. A planning problem is represented by an initial state (objects and propositions about the objects) and a set of goal expressions to achieve. Planning decisions consist of choosing a goal from a set of pending goals, choosing an operator (or inference rule) to achieve a particular goal, choosing a variable binding for a given operator, and deciding whether to commit to a possible plan ordering and to get a new planning state or to continue subgoaling for unachieved goals. Different choices give rise to alternative ways of exploring the search space. By itself, however, PRODIGY cannot interact with other multiagent systems.

Prodigy/Agent [10, 14, 19, 20], written in Allegro Common Lisp 6.2, is a "wrapper" around PRODIGY that allows the automated planner to behave as an independent agent. As such it uses KQML (Knowledge and Query Manipulation Language) [21] as the central agent-communication language and can use agentTool to develop communication protocols. A single copy of Prodigy/Agent can act as a general plan server that may be queried by any heterogeneous agent in a distributed system. Multiple copies of PRODIGY/Agent operate concurrently and coordinate their planning decisions with respect to resource limitations. The most important role for Prodigy/Agent in the MADGS integration is in providing the underlying planning technology for the GTrans interface.

6.2 GTrans

We developed an independent system called GTrans [7, 9, 54, 55] that interacts with the Prodigy/Agent planner using mixed-initiative planning techniques. Using this system, a single human planner can focus on the goals, associated goal priorities, and resource to goal assignments that all three change over time. Rather than thinking of planning as a search mechanism, we present to the user a metaphor of planning as a goal manipulation problem. The primary task is therefore decisions concerning goal change and management. By selecting goal changes, the user can reduce an initial goal to a slightly less demanding goal that partially achieves the state originally sought. Similar selections can also change a goal into a set of distributed subgoals, the achievement of which will

satisfy the super goal. The automated PRODIGY planner provides background support to the user rather than making decisions regarding goal change itself.

From the commander's point of view, planning is achieved through a graphical user-interface that can be manipulated to achieve objectives and project hypothetical situations. As such, it operationalizes an objectives-based planning model [25] and limits the detail thrust upon the human user because the focus of planning is change to the goal rather than the details needed to achieve them.

6.2.1 Goal Transformations

In a dynamically changing environment, aspects that affect a plan and its execution may change at any point. In particular, changes may force are planning effort asynchronously. Traditionally, this determines that a change to the plan be formulated that will allow the goal to be achieved when threatened. However, in many circumstances the goals themselves may need to change rather than the plan *per se* [13]. For example, it makes no sense to continue to pursue the goal of securing an airbase, if the battlespace has shifted to a distant location. At such a point, a robust planner must be able to alter the goal minimally to compensate. Otherwise, a correct plan to secure the old location will not be useful at execution time.

A *goal transformation* represents a goal shift or change. Conceptually it is a change of position for the goal along a set of dimensions defined by some abstraction hyperspace [13]. The hyperspace is associated with two hierarchies. First the theory requires a standard conceptual type-hierarchy within which instances are categorized. Such hierarchies arise in classical planning formalisms. They are used to organize arguments to goal predicates and to place constraints on operator variables. Goal transformation theory also requires a unique second hierarchy.

In normal circumstances the domain engineer creates arbitrary predicates when designing operator definitions. We require that these predicates be explicitly represented in a separate predicate abstraction hierarchy that allows goals to be designated along a varying level of specificity. For example consider the military domain. The domain-specific goal predicate *is-ineffective* takes an aggregate force unit as an argument (e.g., (*is-ineffective enemy-brigade1*)). This predicate may have two children in the goal hierarchy such as *is-isolated* and *is-destroyed*. The achievement of either will then achieve the more general goal [13]. Furthermore if the predicate *is-destroyed* had been chosen to achieve *is-ineffective*, the discovery of non-combatants in the battle area may necessitate a change to *is-isolated* in order to avoid unnecessary casualties. Note also that to defer the decision, the movement may be to the more general *is-ineffective* predicate. Then when the opportunity warrants and further information exists, the goal can be re-expressed. In any case, movement of goals along a dimension may be upward, downward or laterally to siblings.

Goal movement may also be performed by a change of arguments where the arguments exist as objects of or members of the standard object type-hierarchy. The goal represented as the type-generalized predicate (*inside-truck Truck1 PACKAGE*) is more general than the ground literal (*inside-truck Truck1 PackageA*). The former goal is to have some package inside a specific truck (thus existentially quantified), whereas the latter is to have a particular package inside the truck. Furthermore both of these are more specific than (*inside-truck TRUCK PACKAGE*). Yet movement is not fully ordered, because (*inside-truck Truck1 PACKAGE*) is neither more general or less general than (*inside-truck TRUCK PackageA*).

A further way goals can change is to modify an argument representing a value rather than an instance. For example the domain of chess may use the predicate *outcome* that takes an argument from the ordered set of values *checkmate*, *draw*, *lose*. Chess players often opt for a draw according

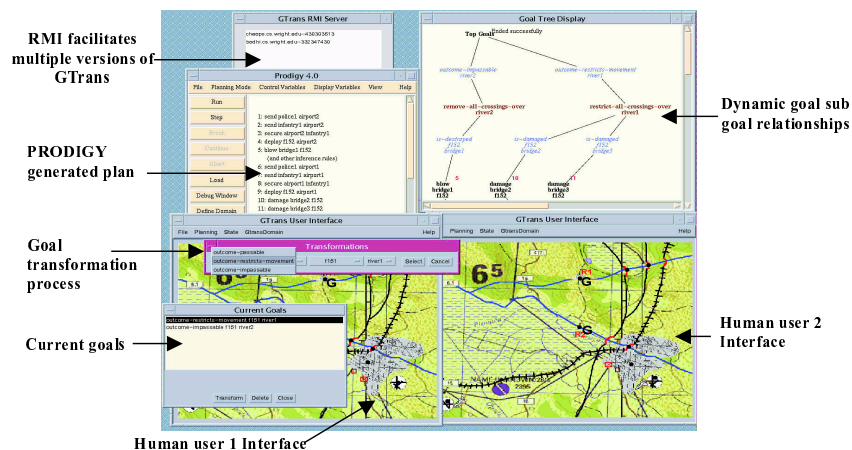


FIG. 6.1. The GTrans interface.

to the game's progress. Thus to achieve the outcome of draw rather than checkmate represents a change of a player's goal given a deteriorating situation in the game.

6.2.2 GTrans User Interface

To directly support the goal manipulation model, we implemented a mixed-initiative interface to a planning system through which the user manipulates goals, the arguments to the goals, and other properties. The interface, written in Java 1.2, hides many of the planning algorithms and knowledge structures from the user and instead emphasizes the goal-manipulation process with a menu-driven and direct manipulation mechanism. GTrans presents a direct manipulation interface to the user that consists of a graphical map with drag and drop capability for objects superimposed upon the map surface. GTrans helps the user create and maintain a problem file that is internally represented as follows. A planning problem consists of an initial state (a set of objects and a set of relations between these objects) and a goal state (set of goals to achieve). The general sequence is (1) to create a planning problem, (2) invoke the underlying planner to generate a plan, and then until satisfied given planning feedback either (3a) change the goals or other aspects of the problem and request a plan or (3b) request a different plan for the same problem. Figure 6.1 shows the interface presented to the user.

We have taken the GTrans system described above and generalized it to allow multiple human planners to operate in teams. Each user has an independently executing copy of GTrans and a supporting component that allows plan manipulation on a common graphical map representation. The systems work independently of each other, yet they have the capacity to take initial conditions and goal descriptions from all or any of the other concurrently running systems. GTrans uses sockets to communicate with Prodigy/Agent and with other executing copies of GTrans. The individual processes may be distributed on remote machines across a network. As such, we are able to fully integrate GTrans into the Carolina environment and provide mixed-initiative planning assistance to the human commander.

7 Putting It All Together

We now describe the prototype system we built in order to test our primary concepts. For our prototype, we focused on demonstrating three capabilities:

- *Mixed-initiative planning and execution.* We achieve this through interactive goal transformation (GTrans) with the human planner and feedback from intelligent resource re-allocation during plan failures. We also provided support to the human user by intelligently re-organizing information that is made available to them.
- *Intelligent logistics assistance.* This is accomplished by assisting the human planner when resource requirements are not met for plan execution. Intelligently and efficiently computing resource re-allocation to provide alternative resources to the human planner while helping mitigate potential conflicts that can arise from multiple missions and goals in the operational theatre.
- *Mobile and multi-agent infrastructure deployment.* We demonstrate that the needs of the two items above can be satisfied in a systematic fashion through knowledge-based software engineering with the automatic generation and deployment of agent components. We focused on communications and mobility in order to efficiently deploy our prototype.

Each of the components we described in Section 2 were integrated and deployed on a mixture of hardware platforms. Our underlying Carolina agent environment was developed in Java and globally executed on Windows, Linux, and Solaris platforms with a mix of Pentium and Sparc CPUs. Our purpose in this mix was to demonstrate the ease of portability and mobility of agents using our approach for cross platform integration. From our results, moving to smaller mobile devices such as PDAs will be readily realizable.

To illustrate the utility of the MADGS system to the commander, we have designed a simple empirical study: Two commanders are operating in the same theatre with independent commands. Each commander is given a mission that potentially conflicts with the other primarily due to resource availability in the theatre. We compose a number of problems where each commander is faced with a number of rivers, each spanned by one or more bridges. The goal is to make each river impassible by destroying the bridges across them. In order to accomplish this task, the commanders with the help of MADGS will assign F-15 tactical fighters from a pool of available aircraft placed at different locations. Each F-15 has the ability to destroy one bridge. However, each commander is only currently aware of the F-15s available near their immediate command. In this military interdiction scenario, commanders must schedule their strikes and will face shortages of F-15s that can only be resolved by communicating and negotiating with the other commander to best utilize all F-15s and achieve their respective missions. An example scenario with rivers and target bridges can be seen in Figure 7.1.

In this section, we first describe our MADGS interface and operations on one such problem scenario above. We then performed an empirical study to determine the overall efficiency and effectiveness of our replanning in the face of resource failure.

Figure 6.1 depicts our GTrans interface that is used by the human commander. For our prototype, we provide a GTrans interface to each user with the added capability of observing other human commanders and their activities to assist in coordination when necessary. This is depicted in the bottom right window in Figure 6.1. As we can see, the mission goals (current goals) are identified and the commander interacts with Prodigy to generate their initial plans with input from the latest information concerning available resources. Once the plan is generated (as seen in window labeled Prodigy 4.0), execution of the plan begins and resources are located to begin scheduling of the plan. Figure 7.3 shows the final confirmation of the plan by the commander together with

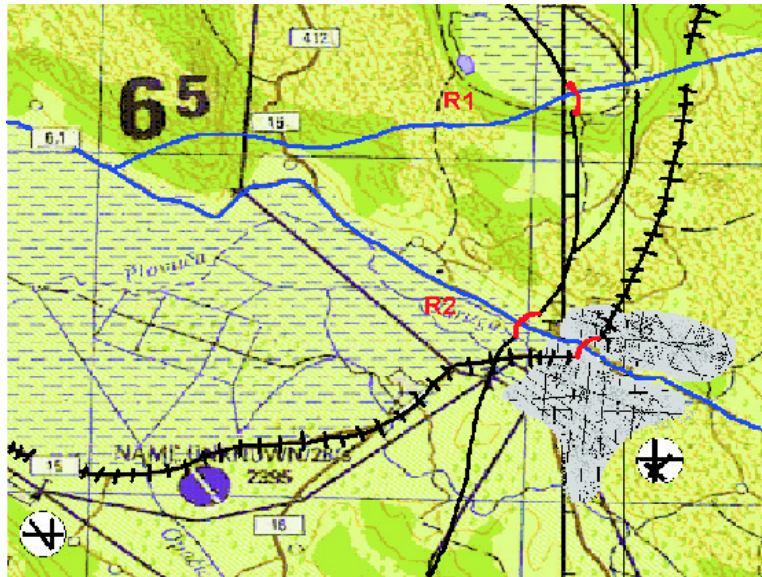


FIG. 7.1. Military Interdiction Scenario. Rivers have target bridges spanning them.

the resource scheduling and execution process. The plan is then sent for execution with agentTool generating new agents in Carolina to carry out the plan execution. Figure 7.2 shows the DGS in action negotiating for resources as needed. Should a plan failure occur at any point, the commander is notified and either resource substitutions are recommended or replanning commences.

We now describe our empirical evaluations of our MADGS system.

7.1 Evaluation

Here, we consider two experiments to gauge the usefulness of MADGS using the scenarios we generated above. Our first experiment considers the usability of our GTrans interface. The second experiment studies the effectiveness of our overall approach to mixed-initiative replanning.

7.1.1 GTrans Evaluation

An experiment was performed with human subjects to compare and contrast the models of search and goal manipulation [9, 54]. The first model is represented by users solving problems in an old interface used in the original PRODIGY planner [13] and the second model is represented by GTrans users. The experiment was designed to evaluate the differences of the two models under differing amount of task complexity using both expert and novices. This experiment uses 18 problems in the military domain as test problems. In these problems, insufficient resources exist with which to solve problems completely. Choices can be made, however, so that a solution is produced that achieves a partial goal satisfaction represented as a ratio of the subject's partial solution to the optimal partial solution.

The graph in Figure 7.4 shows the mean of the goal satisfaction ratio under the goal manipulation model and the search model. When presented with the goal manipulation model, subjects achieve over 95 percent goal satisfaction on average. When presented with the search model, subjects achieve about 80 percent goal satisfaction on average.

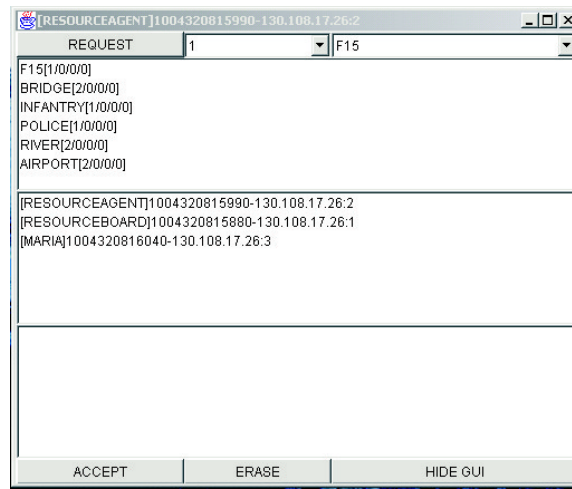


FIG. 7.2. Resource allocation and negotiation.

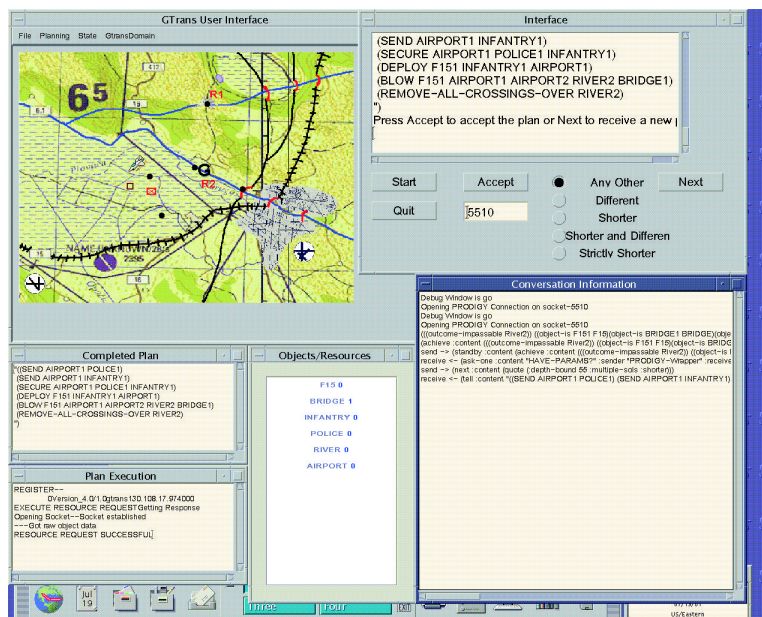


FIG. 7.3. Plan execution commences after initial plan generated.




FIG. 7.4. Goal satisfaction as a function of cognitive model.

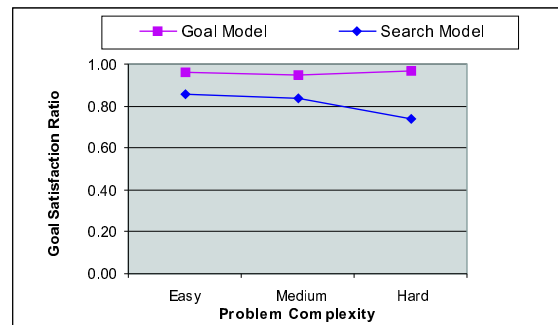


FIG. 7.5. Goal satisfaction as a function of problem complexity.

Given that the cognitive model itself is an important factor as concluded in previous analysis, we next examine the possible relationships among three independent variables: planning model, problem complexity, and subject expertise. Figure 7.5 plots the average goal satisfaction ratio for each combination of the planning model and the problem complexity. As can be observed from the graph, when the goal manipulation model is presented to the user, the goal satisfaction ratio generally remains the same with increasing problem complexity; but when the search model is presented to the user, the goal satisfaction ratio decreases as the problem complexity increases. It is very likely that the effect of the planning model on the user performance depends on the problem complexity.

The next step of our analysis was to examine the possible interaction effects between the planning model and the user expertise level. Figure 7.6 shows the average goal satisfaction ratio for each combination of the planning model and the user expertise level. It is apparent that experts perform better than novices under both planning models. But the two plot lines representing each planning model are not parallel, indicating the possible interactions between the two factors.

7.1.2 Re-planning Study

When most planning systems have to replan given changing environmental conditions such as resource constraints, the system replans completely. That is, if only a single goal fails, the system will

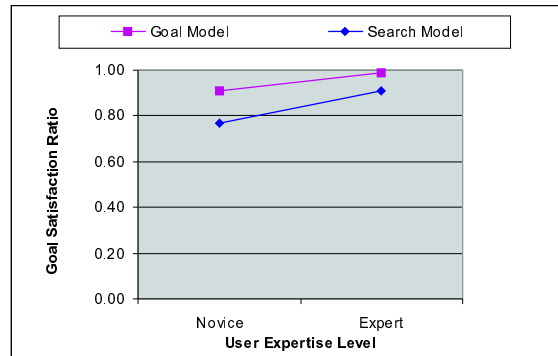


FIG. 7.6. Goal satisfaction as a function of expertise.

still replan for all goals to generate a new solution. Instead, we have extended the Prodigy/Agent component of MADGS to identify the source of the failure and to replan for only that part of the plan that is effected by such failure. We call these two conditions standard replanning and focused replanning.

Our hypotheses is that MADGS will be more efficient in terms measured by the total number of search nodes (planning choice points) expanded and by the total planning time. The first dependent variable measures planning effort directly, because the fewer nodes expanded in the search tree, the fewer poor planning decisions are made. On the other hand time is an indirect measure, because factors other than that spent in planning may effect performance. For example on networked systems, network traffic has an effect.

Using our scenario template above, we generate a number of problems for this experiment. Each problem description contained $2n$ F-15s, where n is the number of rivers in the problem.

In order to test the replanning efficiency of each system (standard vs. focused), 20 tests were run. Each test came from one of four series, containing 2, 4, 6, or 8 rivers. For each series, the number of resource failures was varied from 1 to n . This means that for each test, two plans were computed by the underlying Prodigy/Agent planner within GTrans. The plans are the completely recomputed plan and the revised plan which solves the resource failure only. For example, Test 2.1 has two rivers, and one resource failure. This means that the planner assigns an F-15 to each bridge, and is then notified by Carolina that one of the F-15s assigned is no longer available. This causes the planner to replan for the missing aircraft. Similarly, Test 2.2 has two rivers, but 2 resource failures. The case in which no resource failures occur was not tested since both systems use the same strategy and perform identically in this case. Figures 7.7 and 7.8 depicts the time spent on replanning with and without GTrans and the number of search nodes expanded during the replanning process, respectively.

For the small problems run with these tests, MADGS under the focused condition does not show a significant time savings over the standard condition. In many cases, Tests 2.2, 4.3, 4.4, 6.3 – 6.6, and 8.6 – 8.8, the focus condition requires more time. This can be attributed to two factors: The first is that for small problems like these, the computational overhead of identifying the removable goals is greater than the savings of replanning for a smaller problem. Note that with greater problem complexity, the performance measured by time begins to improve. Future testing should be run to show if this is true by evaluating the two systems on large, real-world problems where the replanning time would be significant. The other factor affecting the times is the planning




FIG. 7.7. Replanning time for standard (without GTrans) and focused (with GTrans).

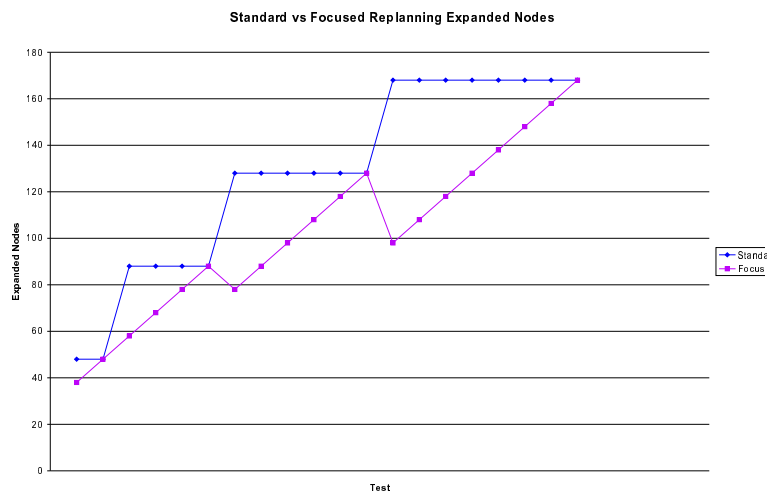


FIG. 7.8. Replanning nodes explored for standard and focused.

environment itself. Each of our tests was run on WSU's lab server. During testing, the machine was being used by other people, causing times to fluctuate depending on the current load. Future testing should be carried out on a separate machine in single user mode in order to isolate the tests from these kinds of problems.

Even with the time overhead issues, the focused condition did have significantly better direct performance over the standard condition when measured by the number of search nodes expanded. Indeed, at no time does the standard condition surpass the focused condition. At most it equals the performance. This is to be expected since the replanning done by the focused condition is only on a subset of the original problems.

8 Conclusion

We presented the MADGS framework for multi-commander dynamic mission planning and execution. We described the major elements of MADGS and illustrated the fundamental logistical and planning support that MADGS can provide to the battlefield commander through our case study. MADGS is currently capable of handling more complicated scenarios which include complex resource substitutions and intelligent allocation recommendations to the commander. Furthermore, in the face of imminent plan failure, MADGS also assists the commander through goal transformations in order to best achieve their given mission. The generation and deployment of agents via agentTool provides a dynamic, efficient, and robust environment that captures the changing nature of battlefield conditions.

One of the elements we intend to pursue next is to address the practical concerns of explaining the choices made (and rejected) by MADGS on recommendations to the commander. Such explanations are doubly critical during on-the-fly planning with partial execution. Hence, we must track the rationale for decisions in order to know when a decision has become inappropriate, due to the changes (either internal or external to the planning system). This will provide the needed transparency of understanding and context to the human commander with regards to MADGS' recommendations.

We must also realize that such complex mission planning will involve disparate types of units/entities; and simply providing a complete global explanation of the entire plan is often counter-productive and counter-intuitive. Most global information in the overall planning is often irrelevant at the lower levels. Thus, the appropriate context must be generated in the explanations with respect to the planning level of the commander(s) we are currently assisting.

The future of assisting the commander on and off the battlefield relies on a foundational understanding of the dynamics of the environment and the requirements of achieving various mission objectives. The MADGS project focuses on addressing the issues of multi-commander, multi-mission planning and execution. In particular, within a single theatre of operation, multiple missions among multiple commanders which are potentially in conflict or competition with one another.

Finally, in the MADGS project, we have also kept an eye towards the current technology trends such as portable mobile computing devices like PDAs, etc. and the potential capabilities that are needed to achieve a MADGS environment. We have envisioned that in the near future, with an understanding of the issues through MADGS, our military commanders, logistics officers, and intelligence analysts can be "armed" with PDAs that provides the critical information they need, when they need it, and also help them in their decision-making. Issues such as incomplete or uncertain information (fog of war) are naturally addressed through a MADGS framework. This paper has detailed the results of our MADGS project and also identified the next steps that should be pursued in reaching our ultimate vision.

References

- [1] Mihai Barbuceanu and Mark S. Fox. Coordinating multiple agents in the supply chain. In *Proceedings of the Fifth Workshops on Enabling Technology for Collaborative Enterprises(WET ICE'96)*, pages 134–141. IEEE Computer Society Press, 1996.
- [2] S. Brown and M. T. Cox. Planning for information visualization in mixed-initiative systems. In M.T. Cox, editor, *Proceedings of the 1999 AAAI-99 Workshop on Mixed-Initiative Intelligence*, pages 2–10, Menlo Park, CA, 1999. AAAI Press.
- [3] Joanna Bryson, Keith Decker, Scott A. DeLoach, Michael Huhns, and Michael Wooldridge. Agent development tools. In *Intelligent Agents VII - Proceedings of the 7th International Workshop on Agent Theories, Architectures, and Languages (ATAL'2000)*, 2000. Springer Lecture Notes in AI, Springer Verlag, Berlin, 2001.
- [4] B. Caldwell. Managing your inventory. *Information WEEK*, 554:88–, 1995.
- [5] J. G. Carbonell, J. Blythe, O. Etzioni, Y. Gil, R. Joseph, D. Kahn, C. Knoblock, S. Minton, A. Perez, S. Reilly, M. M. Veloso, and X. Wang. Prodigy4.0: The manual and tutorial. Technical Report CMU-CS-92-150, Computer Science Department, Carnegie Mellon University, 1992.
- [6] M. T. Cox. A conflict of metaphors: Modeling the planning process. In *Proceedings of 2000 Summer Computer Simulation Conference*, pages 666–671, San Diego, CA, 2000. The Society for Computer Simulation.
- [7] M. T. Cox. A conflict of metaphors: Modeling the planning process. In *Proceedings of the 2000 Summer Computer Simulation Conference*, pages 666–671, 2000.
- [8] M. T. Cox. Interfaces for mixed-initiative planning. In *IUI'2000 Workshop on Using Plans in Intelligent User Interfaces*, Cambridge, MA, 2000. MERL.
- [9] M. T. Cox. Planning as mixed-initiative goal manipulation. In *Proceedings of the Workshop on Mixed-Initiative Intelligent Systems at the 18th International Joint Conference on Artificial Intelligence*, pages 36–41, 2003.
- [10] M. T. Cox, G. Edwin, K. Balasubramanian, and M. Elahi. Multiagent goal transformation and mixed-initiative planning using Prodigy/Agent. In *Proceedings of the 5th World Multi-conference on Systemics, Cybernetics and Informatics, Vol. VII*, pages 1–6, 2001.
- [11] M. T. Cox, B. Kerkez, C. Srinivas, G. Edwin, and W. Archer. Toward agent-based mixed-initiative interfaces. In H.R. Arabnia, editor, *Proceedings of the 2000 International Conference on Artificial Intelligence*, volume 1, pages 309–315. CSREA Press, 2000.
- [12] M. T. Cox and G. Rasul. Human interaction with automated planners through the manipulation and visualization of goal change. Technical Report WSU-CS-00-01, Dept. of Computer Science and Engineering, Wright State University, 2000.
- [13] Michael T. Cox and M. M. Veloso. Goal transformations in continuous planning. In *Proceedings of the 1998 AAAI Fall Symposium on Distributed Continual Planning*, pages 23–30, 1998.
- [14] M.T. Cox, M. Elahi, and K. Cleereman. A distributed planning approach using multiagent goal transformations. In *Proceedings of the 14th Midwest Artificial Intelligence and Cognitive Science Conference*, pages 18–23, 2003.
- [15] Scott A. DeLoach. Analysis and design using mase and agenttool. In *Proceedings of the 12th Midwest Artificial Intelligence and Cognitive Science Conference (MAICS 2001)*, Miami University, Oxford, Ohio, March 31-April 1 2001.
- [16] Scott A. DeLoach, Eric T. Matson, and Yonghua Li. Exploiting agent oriented software engineering in the design of a cooperative robotics search and rescue system. *The International Journal of Pattern Recognition and Artificial Intelligence*, August 2003.

- [17] Scott A. DeLoach and Mark Wood. Developing multiagent systems with agenttool. In *Intelligent Agents VII - Proceedings of the 7th International Workshop on Agent Theories, Architectures, and Languages (ATAL'2000)*, 2000. Springer Lecture Notes in AI, Springer Verlag, Berlin, 2001.
- [18] G. Edwin and M. T. Cox. *Resource coordination in single agent and multiagent systems*, 2001. Submitted.
- [19] M. M. Elahi. A distributed planning approach using multiagent goal transformations. Master's thesis, Wright State University, 2003.
- [20] M. M. Elahi and M. T. Cox. User's manual for Prodigy/Agent, Ver 1.0. Technical Report WSU-CS-03-02, Dept. of Computer Science and Engineering, Wright State University, 2003.
- [21] T. Finin, D. McKay, and R. Fritzson. An overview of KQML: A knowledge query and manipulation language. Technical report, Computer Science Department, University of Maryland, 1992.
- [22] Mark S. Fox, John F. Chionglo, and Mihai Barbuceanu. The integrated supply chain management system. Technical report, University of Toronto, 1993.
- [23] Nobuhiro Kataoka, Hisao Koizumi, and Hiedeaki Simizu. Architecture of an autonomous distributed system and verification of implementation as a logistics information management system. In *Proceedings of the 1997 3rd International Workshop on Object-Oriented Real-Time Dependable Systems (WORDS'97)*, 1997.
- [24] Elizabeth Kendall. Agent roles and role models: New abstractions for multiagent system analysis and design. Proceedings of the International Workshop on Intelligent Agents in Information and Process Management, Bremen, Germany, September 1998, September 1988.
- [25] G. A. Kent and W. E. Simons. Objective-based planning. In P. K. Davis, editor, *New Challenges for Defense Planning: Rethinking How Much is Enough*, pages 59–71. RAND, 1994.
- [26] B. Kerkez and M. T. Cox. Planning for the user interface: Window characteristics. In *Proceedings of the 11th Midwest Artificial Intelligence and Cognitive Science Conference*, pages 79–84, Menlo Park, MA, 2000. AAAI Press.
- [27] B. Kerkez, M. T. Cox, and C. Srinivas. Planning for the user interface: Window content. In H.R. Arabnia, editor, *Proceedings of the 2000 International Conference on Artificial Intelligence*, volume 1, pages 345–351. CSREA Press, 2000.
- [28] Kwindla Hultman Kramer, Nelson Minar, and Pattie Maes. Tutorial: Mobile software agents for dynamic routing. <http://www.media.mit.edu/nelson/research/routes/sigmobile.ps>, 1999.
- [29] E. Kutanoglu and S.D. Wu. *An Auction-Theoretic Modeling of Production Scheduling to Achieve Distributed Decision Making*. PhD thesis, Dept. of Industrial and Manufacturing Systems Engineering, Lehigh University, 1999.
- [30] Timothy H. Lacey. A formal methodology and technique for verifying communication protocols in a multi-agent environment. Afit/eng/00m-12, School of Engineering, Air Force Institute of Technology (AU), Wright-Patterson Air Force Base Ohio, USA, 2000.
- [31] Timothy H. Lacey and Scott A. DeLoach. Automatic verification of multiagent conversations. In *Proceedings of the 11th Annual Midwest Artificial Intelligence and Cognitive Science Conference*, Fayetteville, Arkansas, April 2000.
- [32] Jyi-Shane Liu and Katia P. Sycara. Coordination of multiple agents for production management. *Annals of Operations Research*, 75:235–289, 1997.
- [33] Peter B. Luh and Debra J. Hoitomt. Scheduling of manufacturing systems using the lagrangian relaxation technique. *IEEE Transactions on Automatic Control*, 38:1066–1080, 1993.

- [34] J. Todd McDonald, Michael L. Talbert, and Scott A. DeLoach. Heterogeneous database integration using agent oriented information systems. In *Proceedings of the International Conference on Artificial Intelligence (IC-AI'2000)*, Monte Carlo Resort, Las Vegas, Nevada, June 26-29 2000.
- [35] Scott A. O'Malley, Athie L. Self, and Scott A. DeLoach. Comparing performance of static versus mobile multiagent systems. In *National Aerospace and Electronics Conference (NAECON)*, Dayton, OH, October 10-12 2000.
- [36] Philippe Quindec and Grard Padiou. Flight plan management in a distributed air traffic control system. In *First International Symposium on Autonomous Decentralized Systems (ISADS-93)*, 1993.
- [37] Marc J. Raphael and Scott A. DeLoach. A knowledge base for knowledge-based multiagent system construction. In *National Aerospace and Electronics Conference (NAECON)*, Dayton, OH, October 10-12 2000.
- [38] David J. Robinson. A component-based approach to agent specification. Afit/gcs/eng/00m-22, School of Engineering and Management, Air Force Institute of Technology (AU), Wright-Patterson Air Force Base Ohio, USA, March 2000.
- [39] G. Mitchell Saba and Eugene Santos, Jr. The multi-agent distributed goal satisfaction system. In *Proceedings of the International ICSC Symposium on Multi-Agents and Mobile Agents in Virtual Organizations and E-Commerce*, pages 389–394, 2000.
- [40] Norman M. Sadeh, David W. Hildum, Dag Kjenstad, and Allen Tseng. Mascot: An agent-based architecture for coordinated mixed-initiative supply chain planning and scheduling. In *Proc. 3rd Int'l. Conf. on Autonomous Agents (Agents'99) Workshop on Agent-Based Decision Support for Managing the Internet-Enabled Supply Chain*, Seattle, WA, 1999.
- [41] Eugene Santos, Jr., Feng Zhang, and Peter B. Luh. Multi-agent logistics management. In *Proceedings of the International Conference on Internet Computing (IC '2001)*, pages 240–246, 2001.
- [42] Eugene Santos, Jr., Feng Zhang, and Peter B. Luh. Intra-organizational logistics management through multi-agent systems. *Electronic Commerce Research*, 3:337–364, 2003.
- [43] A. Self. Design and specification of dynamic, mobile, and reconfigurable multiagent systems. Afit/eng/01m-11, School of Engineering, Air Force Institute of Technology (AU), Wright-Patterson Air Force Base Ohio, USA, March 2001.
- [44] Weiming Shen and Douglas H. Norrie. An agent-based approach for manufacturing enterprise integration and supply chain management. In G. Jacucci, editor, *Globalization of Manufacturing in the Digital Communications Era of the 21st Century: Innovation, Agility, and the Virtual Enterprise*, pages 579–590. Kluwer Academic Publishers, 1998.
- [45] Leyuan Shi, Chun-Hung Chen, and Enver Ycesan. Simultaneous simulation experiments and nested partition for discrete resource allocation in supply chain management. In *The 1999 Winter Simulation Conference (WSC'99)*, 1999.
- [46] C. H. Sparkman. Transforming analysis models into design models for the multiagent systems engineering methodology. Afit/eng/01m-21, School of Engineering, Air Force Institute of Technology (AU), Wright-Patterson Air Force Base Ohio, USA, March 2001.
- [47] C. Srinivas, M. T. Cox, and V. Laxminarayanan. Gtrans user manual and reference. Technical Report WSU-CS-00-02, Dept. of Computer Science and Engineering, Wright State University, 2000.
- [48] M. M. Veloso, J. G. Carbonell, A. Perez, D. Borrago, E. Fink, and J. Blythe. Integrating planning and learning: The PRODIGY architecture. *Journal of Experimental and Theoretical Artificial Intelligence*, 7(1):81–120, 1995.

- [49] T. Wagner, A. Garvey, and V. R. Lesser. Design-to-criteria scheduling: Managing complexity through goal-directed satisficing. In *Proceedings of the AAAI Workshop on Building Resource-Bounded Reasoning Systems*, 1997.
- [50] William E. Walsh and Michael P. Wellman. A market protocol for decentralized task allocation. In *Proceedings of the Third International Conference on Multi-Agent Systems*, pages 325–332, Paris, France, 1998.
- [51] M.P. Wellman. Market-oriented programming: Some early lessons. In S. Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*, chapter 4. World Scientific, 1996.
- [52] Stanley K. Yung and Christopher C. Yang. A new approach to solve supply chain management problem by integrating multi-agent technology and constraint network. In *Proceedings of the 32nd Hawaii International Conference on System Sciences(HICSS-1999)*, Maui, Hawaii, 1999.
- [53] Daniel Dajun Zeng and Katia Sycara. Dynamic supply chain structuring for electronic commerce among agents. In Matthias Klusch, editor, *Intelligent Information Agents*. Springer, 1999.
- [54] C. Zhang. Cognitive models for mixed-initiative planning. Master’s thesis, Wright State University, 2002.
- [55] C. Zhang, M. T. Cox, and T. Immaneni. GTrans version 2.1 user manual and reference. Technical Report WSU-CS-02-02, Dept. of Computer Science and Engineering, Wright State University, 2002.