REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
14-March-2006

2. REPORT TYPE
final report

3. DATES COVERED *(From - To)*
28-Nov-2006 28-March-2007

4. TITLE AND SUBTITLE
Distributed Planning and Coordination of Team-oriented Activities

5a. CONTRACT NUMBER
N62558-06-P-0353

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Michal Pechoucek, Austin Tate, Eduard Semsch, Jeff Dalton, Gerhard Wickler

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Gerstner Laboratory, Czech Technical University in Prague
Technicka 2, 166 27, Prague 6, Czech Republic
Artificial Intelligence Application Institute
School of Informatics, University of Edinburgh

8. PERFORMING ORGANIZATION REPORT NUMBER
1

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

13. SUPPLEMENTARY NOTES

14. ABSTRACT
During the second phase of the N62558-06-P-0353 project the infrastructure for joint research work among ATG and AIAI has been established, interoperability of the HTN/I-X technology with a/globe multi-agent system has been achieved, more complex experiment has been performed and the proposal for the joint research project has been elaborated and submitted. This report a review of the state of the art of the relevant work, basic description of the integrated architecture and three separate demonstration scenarios. The Research reported in this document has been made possible through the support and sponsorship of the U.S. Government through its European Research Office of the U.S. Army. This report is intended only for the internal management use of the Contractor and U.S. Government.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Michal Pechoucek
				31	**19b. TELEPHONE NUMBER** *(Include area code)* +420 723 4949456

Distributed planning in a dynamic, semi-trusted and opportunistic environment

Michal Pechoucek[1], Gerhard Wickler[2], Austin Tate[2], Eduard Semsch[1] and Jeff Dalton[2]
[1]Czech Technical University, Department of Cybernetics, Agent Technology Group
[2]University of Edinburgh, School of Informatics, Artificial Intelligence Applications Institute

Abstract: During the second phase of the N62558-06-P-0353 project the infrastructure for joint research work among ATG and AIAI has been established, interoperability of the HTN/I-X technology with a/globe multi-agent system has been achieved, more complex experiment has been performed and the proposal for the joint research project has been elaborated and submitted. This report a review of the state of the art of the relevant work, basic description of the integrated architecture and three separate demonstration scenarios.
The Research reported in this document has been made possible through the support and sponsorship of the U.S. Government through its European Research Office of the U.S. Army. This report is intended only for the internal management use of the Contractor and U.S. Government.

1 Introduction

In this paper we will discus the planning problem positioned in a very specific environment formulated in parts by the project funding agency[1]. The environment:

- is to be **non-centralized** and with flat organization hierarchy [R1] – the existence of a central coordinating and planning process shall be brought to absolute minimum and the planing knowledge, information about ac-

[1]ERO - European Research Office of US Army

tors skills, resource availability knowledge and goals perception shall be distributed,

- shall provide **partial knowledge sharing** [R2] – the actors in the environment are motivated to keep substantial part of their private planning knowledge and resource availability information undisclosed,
- shall allow **varying interaction availability** [R3] – based on communication infrastructure featuring partial and temporal inaccessibility due to e.g. ad-hoc networking, unreliability of the communication infrastructure or actors to change off-line/on-line status,
- is to be **very dynamic** [R4] where both resource availability as much as goals persistence is expect to be changing between the planning and execution phase, while also during the execution phase, and
- is to be **opportunistic** [R5] – allowing the actors reason about potential goal accomplishment opportunities that may arise in the environment and also consider opportunities of the collaborating actors in the environment.

Such a set of requirements is typical for rescue operations, complex humanitarian missions, other OOTW large scale multi-national coalition operations as well as small size military combat ops. Such features are also typical for complete different set of application domains such as virtual organizations and social networking.

1.1 An Abstract Deployment Scenario

The targeted deployment scenario is going to cover multiple levels of the planning and execution process within a dynamically developing situation, involving a number of:

- humans (command authorities, planners and operators),
- mobile vehicles (such as trucks and multiple UAVs with appropriate sensor suites),
- a network of unattended ground sensors (UGS),
- software agents (planning systems, sensor processing systems, semantic web resources of various kinds, etc.).

Each of these actors will be modeled as either an autonomous agent or a container hosting several intelligent agents. An example of a container could be a UAV bird hosting various software agents implementing the various aspects of its autonomous aspects (such as its planning agent, sensor processing agent, communication and negotiation agent, etc.).

The to-be-developed multi-agent system will initiate the process of negotiation between humans, autonomous robotic agents, and virtual agents to acquire and maintain geographical and contextual information and to exchange plan-related information in a timely fashion.

The final scenario will demonstrate automatic re-tasking of the autonomous agents (e.g. UAVs and other vehicles) through an agent-oriented collaborative process, where the devices with the appropriate capability (e.g., IR sensors, still or streaming imagery, chemical sensors) could be re-tasked based on some mission priority. When re-tasking, other capabilities might need to be reconfigured autonomously to ensure that they can maintain their overall mission tasks in the altered context (e.g., providing communications connectivity over a large geographic area for a mobile force).

The final scenario shall demonstrate that the cognitive load on the warfighter is not increased and its ease of use and transparency is maintained. The evaluation may require access to end users to show these properties.

1.2 Planning Levels

The distritbuted planning and coordination task will be addressed using a hierarchy of planning levels that correspond to very different perspectives:

1. **strategic** - problem analysis, sense-making and high-level strategic task setting and approach selection. For this level we plan to use Compendium (http://www.compendiuminstitute.org) as a basis since this has already been demonstrated and successfully evaluated as useful in support of human military planners. There is the potential to use an issue-based approach to sense-making, option analysis, argumentation and decision support at this level. A combined Compendium and I-Plan tool was demonstrated and evaluated during the Collaborative Operations for Personnel Recovery (Co-OPR) project as part of Experiment B of DARPA's Integrated Battle Command program [Tate and Selvin, 2006][2].

2. **operational** - generation of responses/plans and refinement/repair of these dynamically as needed. This level could use mixed-initiative multi-agent planning approaches and a HTN approach as these have been proven to be at a level that relates well to human planners, who need to maintain and communicate a plan at a suitable level of abstraction. HTN approaches and plan representations also act as a bridge between the levels. This level could be based on an I-X Process Panel [Tate *et al.*, 1999; Tate and Stader, 2002; Wickler, 2006]. Its HTN planner, I-Plan, will need to be developed further. Refinement of the algorithms to add temporal and consumable/renewable resource constraint management at least as included in the earlier O-Plan planner [Currie and Tate, 1991a] will be necessary [Drabble and Tate, 1994]. A simple spatial constraint manager and/or specialised sub-solvers for route planning could be added, as allowed for, but not yet explored, in its distributed planning architecture [Tate *et al.*, 1994]. Multi-level plan execution monitoring [Reece and Tate, 1994] and repair algorithms already available for O-Plan should also be

[2]See also http://www.aiai.ed.ac.uk/project/co-opr/expt/

incorporated. The aim is to bring I-Plan as a module up to a level of competence to act as a restartable/incremental planner able to refine multiple options for a response and change these as circumstances change.

3. **tactical** - adjustment and further more detailed refinement of the plans to meet the stated objectives within the constraints given, but talking into account local circumstances and context. Here agent-based techniques [Rehak *et al.*, 2006; Pechoucek *et al.*, 2006] will be used for peer-to-peer negotiation among individual actors aimed at optimal responsibility delegation and resource allocation. ACROSS and **a/globe** [Šišlák *et al.*, 2006] are anticipated as providing the basis for this tactical planning layer, though we anticipate that we will make available a version of the I-Plan planning service to the tactical agents to assist them in refining and adapting their plans to local circumstances, as was done in the O-Plan/WorldSim Operational to Tactical planning and plan execution support techniques [Reece and Tate, 1994; Tate, 1984; Tate, 1989]. Reporting to higher levels should be done, and this may include some details of the approach taken (especially if this raises new issues, or introduces other mutually constraining information), progress and completion (whether successful, partially successful or failed).

The levels will be tied together through the use of a shared underlying plan representation based on `<I-N-C-A>` [Tate, 1993; Tate, 2000a; Tate, 2003] which can act as an "intelligible" bridge between such levels [Tate, 1996; Tate, 2000b; Siebra and Tate, 2006].

2 Relevant Work

2.1 Extensions to the Classical Action Representation

In this section we will describe some extensions to the classical planning representations that would benefit the envisaged scenario described above. While these mostly do not directly address the objectives set out at the beginning of this document, they are relevant for some more fundamental issues and the scenario we intend to develop.

2.1.1 Relaxing Classical Assumptions

Classical representations for the planning problem often make a number of assumptions [Georgeff, 1987; Ghallab *et al.*, 2004] that are unrealistic in many domains. Here we will briefly discuss the assumptions we believe to be unrealistic for the envisaged scenario and a wider, military context. Some of the approaches that have been used to avoid these assumptions will also be discussed.

The assumption that there is only a finite number of world states is unrealistic in domains where we can create new objects or have to deal with numeric parameters. Both is clearly the case in our scenario where new resources can be produced and available quantities of resources need to be known. This means that planning algorithms like Grahplan or SAT-based techniques will be difficult to apply here. Classical approaches can be extended using functions to deal with object creation [Ghallab *et al.*, 2004, parts I and III], and resources are usually dealt with by schedulers as discussed in section 2.2.2 below.

The assumption that the world is fully observable and complete knowledge of the state of the world is available is unrealistic as there may be an adversary with unknown position and intensions. General techniques for dealing with partially observable states and planning under uncertainty are described in [Ghallab *et al.*, 2004, part V]. One approach to dealing with partially observable worlds is to provide the planner with explicit advice that guides an exhaustive search [Kearns *et al.*, 2000] or, to use machine learning techniques to get at such knowledge [Even-Dar *et al.*, 2005]. Another approach to addressing the problem of uncertainty is to make some reasonable assumptions about the world which are known in many domains. In [Albore and Bertoli, 2006] an approach using such assumptions to constrain the search is described. They encode the assumptions in linear temporal logic (LTL) and describe a forward-chaining algorithm that generates a conditional, assumption-based plan. Since the assumptions may not hold, this approach shifts some of the problem to the execution phase where the assumptions need to be monitored and replanning may occur (see section 2.3.4 below). Taking the idea of shifting effort to execution time further, [Chang and Amir, 2006] describes an planning algorithm that interleaves planning and execution, planning ahead only as makes sense given the uncertainty, then executing, etc. For a limited class of actions (deterministic, without conditional effects), the algorithm is guaranteed to reach a goal state in a bounded number of steps.

The assumption that all actions are deterministic often quite reasonable, but certain actions tend to have less predictable outcomes. If an adversary is involved, actions can appear to be non-deterministic, and this will be the case for our scenario. This kind of domain usually calls for plan that may contain branching actions where the branching condition is an observation action, but this requires an observable world. If there is no observability this problem is called conformant planning, and this has recently been shown to be in 2EXPSPACE [Bonet, 2006], depending on the type of states represented. In reality, this means that there is little hope that such a problem can be solved in a scalable way in the near future. Even with full observability the complexity results are not encouraging [Smith and Weld, 1998]. Despite the theoretical complexity there has been recent progress, however. In [Bryce *et al.*, 2006] a heuristic search algorithm is described that appears to scale better than currently dominant CSP/SAT-based approaches. The heuristic used is an adaptation of a planning graph heuristic, and the resulting probability distributions are computed using Sequential Monte Carlo methods [Doucet *et al.*, 2001].

The assumption that there are no events that may occur is of course unrealistic and not applicable in the scenario we envisage. For example, changes to the weather condition are events. Effectively, events can be treated as adding a non-deterministic component to actions though, which means that they present no further complication in computational terms.

The assumption that a goal must be stated as a set of explicit goal conditions is unrealistic in the sense that there is often more to a goal than just this. We will discuss this issue further in the section 2.1.2 as this calls for different types of representations of planning problems. Similarly, we will come back to the assumption that time only has to be implicit in the plan in section 2.2.1, where we will discuss constraint mangers that can deal with time.

Finally, the assumption that planning takes place offline is only realistic if the planning is fast compared to the actions and events that are in the planning domain. This will usually be the case for the level of planning that we envisage here, e.g. the kind of planning done in a command and control center. Should this turn out to be too coarse grained and online planning is required, an overview of some techniques to deal with this problem can be found in [Ghallab *et al.*, 2004, part VI].

There are several further assumptions that the are not valid in the investigated type of scenarios. These assumptions have not been explicitly listed in [Ghallab *et al.*, 2004], but they are assumed to be valid for all the centralized planners.

The assumption that the planner works on top of the centralized planning knowledge and data about type and availability of resources is not realistic in our scenario. By distribution of the planning knowledge we understand the situation where the set of all the available operators (or refinements in hierarchical planning) is not known by the planner. Therefore, it may not possible for the planner to find a complete and consistent plan.

The assumption, very often made in the domain of multi-agent systems, that all the data and knowledge can be communicated among the individual actors is not realistic in our scenario either. Firstly, in complex and dynamic environment it can never be assured that all the actors provide all the data and knowledge in real time and that communicated data will be of a reasonable size. Additionally, the actors may not be willing to provide all the available upon request due to confidentiality reasons.

Data may not be as easy to be communicated due to the fact that some of the actors may become inaccessible in our domain scenarios. Such a disconnection may be caused be infrastructure failure, actor becoming intentionally off-line or properties in an ad-hoc network.

Often it is assumed that the planning can be put separate from the resource allocation phase. In the situation of our scenario we need to produce plans that are possible to be implemented in the distributed settings. As resource availability is dynamically changing in our environment we need to integrate planning and resource allocation (that we will refer to as task delegation in the setting of distributed planning). In [Ghallab *et al.*, 2004, chapter 15] there is presented a mechanism for integrating planning and resource allocation that is based on integration of causal and resource oriented reasoning. However these methods cannot be used due to violation of the assumption of knowledge and data availability from the point of the planner.

The result of classical planning process is a plan that is complete and consistent. Therefore plan that can be achieved. When planning and task delegation is integrated we seek a plan that is not only doable in theory but can be also achieved with respect to available resources. In our domain we however seek problems that will be implemented. Consequently plans, for which the actors get committed. Reasoning about the level of commitments is an important aspect of planning in distributed, highly dynamic environment.

2.1.2 Extended Representation

Conditional effects are used to model actions in which the outcome of performing the action depends on the situation in which the action was performed, e.g. toggling a light switch.

Universally quantified affects can be seen as a more general version of conditional effects, where an effect applies all objects that satisfy certain conditions. For example, when a truck drives somewhere, everything that is loaded on the truck will change location, too. Since we are dealing with this kind of transportation in our scenario, this would seem like a useful extension. In finite domains this can done relatively easily as the universal quantification usually only applies to a finite set of objects that can be enumerated. For an infinite domain the problem has not been addressed as far as we are aware. This extension is included in the action description language (ADL) [Pednault, 1989; Pednault, 1994]

An operator with disjunctive preconditions can usually be decomposed into multiple operators with each disjunct as a preconditions. Thus, this extension requires only a little pre-processing. Whether it is desirable to have this extension in our scenario is not clear.

Axiomatic inference is used to reason about the state of the world beyond what is asserted by the effects of the actions in the (and implicit frame axioms). This can be very useful for automating certain types of analysis or abstracting away from the current situation for sense-making purposes. Early work on AI planning was, of course, based on axiomatic inference [Green, 1969] which would make it easy to integrate. A different approach distinguishes primary predicates for domain relations that can be preconditions and effects of operators, and secondary predicates that are generated by axiomatic inference and may only appear in preconditions, but not effects of operators. As such, axiomatic inference has been used in several planners [Ghallab *et al.*, 2004, chapters 10 and 11].

Introducing functions into the basic state representation significantly increases the expressive power of the representation, but it means that reasoning about single states becomes an undecidable problem, which raises severe complexity issues [Ghallab *et al.*, 2004, chapter 3]. Only strictly limited use of functions may be of use in practice. The problem is that anything to do with computation and numbers almost certainly requires functions. For example, the range of a helicopter given its current fuel level requires functional representations.

One way of addressing this problem are attached procedures, that is, code that is attached to a function symbol and gets evaluated at the right time in the planning procedure. Similarly, there can be code attached to predicates to evaluate them in a given situation. While this is a practical solution, it defies any theoretical claims about the planning procedure, including its correctness which is usually a minimal requirement.

Extended goals allow for a problem specification that goes beyond a simple set of conjunctive goal conditions, e.g. states to be avoided or utility functions that make planning an optimization problem. Clearly such extensions could be very useful in our scenario, although it is expected that the human controller using the planing tool would take the final decision when it comes to utility. Some types of extended goals and approaches to dealing with them are described in [Ghallab *et al.*, 2004, part V]. A variant of the classical planning problem is over-subscription planning (OSP) [Smith, 2004] where it is no longer assumed that all goal conditions can be achieved simultaneously or with the given resources. Actually, there are different formulations of this problem [Brafman and Chernyavsky, 2005; Sanchez and Kambhampati, 2005]. Recently, Stochastic OSP, which allows for uncertainty, has been addressed using tools developed for hierarchical reinforcement planning to exploit a hierarchical structure in such problems [Meuleau *et al.*, 2006]. Instead of solving one large MDP, their algorithm solves a number of smaller MDP which can be done more efficiently.

A number of these extensions (or approximations thereof) have been included in the Planning Domain Definition Language (PDDL) [McDermott, 1998; Fox and Long, 2002] used in the international planning competition. A number of planners exist that can read domain definitions in PDDL, but a specification for more practical HTN-style refinements was only included in version 1, reflecting the rather academic interest in the planning competition.

From the viewpoint of distributed planning, the planning problem representation needs to be further extended so that the unsatisfied operators and/or undecomposed refinements that form an incomplete plan can be further investigated. The incomplete plan can be made complete by requesting collaborating actors for knowledge about the unsatisfied operators. E.g. in hierarchical planning it needs to be made possible to work with refinements that are neither terminal nor for them an additional lower level refinement exists. In such a situation the request for the refinement rule would be broadcasted among the planning actors.

Knowledge collection, required during such planning process can be implemented by a specific type of procedural attachments, that would initiate *contract-net-protocol*-like [Smith, 1980], asynchronous interaction process. Asynchronicity and uncertain outcome of such interaction is a key bottleneck of integration of planning and negotiation processes. It is impossible to reason about such plans effectively a'priori. The same applies for the process of resource allocation integrated within the planning process (which is going to be the case of our domain scenario).

Another alternative of how to work with partial, incomplete plans is to support integration of several, ad-hoc created partial plans by means of *Multi-Agent Plan Coordination* methods.

An important extension of the current plan representation in the domain of distributed planning is oriented towards enforcement and monitoring of the collaborative action.

2.2 Constraint Managers

In this section we will look at some of the types of constraints that need to be managed with plans in the envisaged scenario.

2.2.1 Time Constraints

A quite concise overview of the representation of and reasoning about time in the context of planning can be found in [Ghallab *et al.*, 2004, sections 13.5 and 14.4]. A lot of this work can be considered well established now and efficient algorithms that scale well are known. The major HTN planning systems that are currently in use in a number of domains [Currie and Tate, 1991b; Wilkins, 1988; Nau *et al.*, 2001] all include a temporal constraint manager that handles either time point or interval constraints.

Another HTN planner that includes a temporal constraint manger is SIADEX, and some performance improvements have been achieved here recently [Castillo *et al.*, 2006; Fdez-Olivares *et al.*, 2006]. This planner records the causal rationale behind primitive actions, a common technique in partial-order planners, and exploits this structure when it comes to constraint propagation in the temporal network. This algorithm is based on the well-known P-C2 algorithm [Dechter, 2003]. This work is not only interesting from a technical perspective though. Its application domain is in the area of emergency response (forrest fire fighting) and thus quite closely related to the kind of scenario we envisage.

2.2.2 Resource Constraints

Again, a quite concise overview of the representation of and reasoning about resources (scheduling) in the context of planning can be found in [Ghallab *et al.*, 2004, section 15.5]. The scheduling problem itself has been addressed in many years of active research and efficient algorithms are known for many variants of the problem. What we are interested here is the integration of scheduling into planning algorithms and there are a number of recent results that are worth a second look.

Supply chains are ubiquitous not only in manufacturing, but also play an important role in military operations and emergency response tasks. Traditionally, supply chains are not an automated process but recent advances in planning and scheduling technology are moving towards this [Chen *et al.*, 1999; Dadeh *et al.*, 1999]. The Trading Agent Competition for Supply Chain Management (TAC SCM) has proved a good testbed for systems addressing this problem and the possibility to evaluate new ideas in a competitive setting has led to an increase in research in this area. Recently, TacTex-05 [Pardoe and Stone, 2006], the winning agent from the 2005 competition, used a predictive approach for future resource availability and constraints, a technique that has been used in supply chain management for a long time, and integrated this into the planning and scheduling system.

A problem similar to the Stochastic OSP problem described in section 2.1.2 above, a variant of the scheduling problem known as Optimal Competitive Scheduling deals with the problem of finding a schedule by packaging as many jobs as possible into a given time horizon, given a limited and usually insufficient set of resources for all jobs. The aim is to find a schedule that is optimal according to some utility function. While this problem is NP-complete in general, a number of tractable cases have been identified including problems where activities have fixed start times, duration and value [Sandholm and Suri, 2000], problems with resolved resource conflicts and some restrictions on preference functions [Morris *et al.*, 2004], and problems with limitations on the objective function or which can be reformulated as a valued CSP [Frank *et al.*, 2006].

In a dynamic world it is usually not feasible with a complete plan or schedule that will solve a given problem. Instead, planning and scheduling have to take place incrementally as new problems and opportunities arise. While high-level plans can remain relatively stable, low-level schedules tend to require more of a dynamic approach, and a number of approaches to incremental scheduling have been described in the literature [Smith, 1994; El Sakkout and Wallace, 2000; Gallager *et al.*, 2006].

2.3 Dynamic Multi-Agent Planning Techniques

In this section we will look at the planning techniques and algorithms that can be used to address specific sub-problems that may occur in the envisaged scenario.

2.3.1 Multi-Agent Plan Coordination

The problem of coordination of a number of partial plans elaborated by a collective of self-interested actors is supported by several different techniques. The classical work of e.g. *Multi-Agent Plan Coordination* (MAPC) [Cox and Durfee, 2005] is based on partial order causal link definition of a plan, provides a formal definition of the multi-agent parallel partial order causal link plan by introducing *parallel step thread flaw* and *plan merge step flaw*. Multi-agent plan coordination process is implemented by means of branch-and-bounds-like search through the space of step thread and merge step flaws. This highly relevant, formally well founded work provides empirical comparison to classical work [Yang, 1997]. The difficulty with this approach is that it is fully centralized. Decentralization of plan merging proces would be inevitable in our domain scenario.

A fully distributed plan merging approach is presented in [van der Krogt *et al.*, 2003], where the autonomous actors are merging their plans by means of sharing resources. This is an optimization oriented approach, where each agent is trying to maximize its profit from reducing its plan and merging it with other agents' plans. Plan merging is implemented by one step-wise auctioning procedure, where the agents are negotiating the most profitable release of a specific action and thus sharing the resources. The presented approach rely on a a trusted third party auctioneer, which is unrealistic in our domain. The algorithm is not guaranteed to find an a global optimum, as the auctioneer is choosing the most cost reducing resource release each round.

Besides the two above listed approaches to post-planning coordination, there are works of [von Martial, 1992; Tonino *et al.*, 2002], that is aimed at revisions of the agents' individual plans so that conflicts are resolved and resources are shared. of the

As stated in [Witteveen and de Weerdt, 2006] the plan coordination problem among self-interested, semi-trusted and noncooperative agents needs to be often achieved by preplanning coordination, that imposes a minimal set of additional constrains to the original planning problem. Examples of such methods are e.g. portioning method [Valk *et al.*, 2005] or temporal decoupling [Hunsberger, 2002].

Multi-agent opportunistic planning is a very specific technique for collaborative planning and collaborative plan execution with the aim to utilize the best sharing resources and sharing overlapping goals. The key idea is that each agent creates plans that also include opportunities for the other agents. If the opportunity goal becomes pending it can be achieved by other agents. Goals may became unachievable due to changes in the environment or were unachieveble from the very start. They work listed in [Lawton and Domshlak, 2004] is based on use of *partial order planning graphs* (POPG) as they provide a bigger deal of flexibility for execution. In [Lawton and Domshlak, 2004] they provide tests of several different strategies for selecting the additional goals for which an opportunity may arise. Multi-agent opportunistic planning is based on minimal knowledge sharing (they share information about other agents capabilities and assigned goals, which may be even too much in our domain). The current implementation does not allow for online replanning (in a sense of dropping plans and adopting new plans instead).

2.3.2 Task Delegation

No matter if the plans are coordinated by pre-planning or post-planning methods, it needs to be optimally decided which subtasks are allocated to which agents. Under the assumption of decentralized availability of information about individual agents' skills and availability, the multi-agent community provides a long list of contract-net-proposal [Smith, 1980] like methods and variants of combinatorial auctioning methods [Hunsberger and Grosz, 2000; Boutilier *et al.*, 1999; Sandholm, 2002].

There are two key challenges that our domain implies: (*i*) avoiding local optimum of the task allocation in combinatorial auction problems (as they do not allow backtracking) and (*ii*) coping with agents' reluctance to provide the complete information about their resource availability. The latter is the case of the semi-trusted communities mainly [Pěchouček *et al.*, 2002].

While it is very difficult (nearly to impossible) to establish global social welfare optimal allocation of tasks in the dynamic environment, there are several methods that extend classical contract-net-protocol in order to allow optimization across a small number of separate negotiation rounds[3]. *Extended Contract Net Protocol* (ECNP) [Fischer *et al.*, 1995; Fung and Chen, 2005] extends the protocol with a `provisional accepts` and `provi- sional rejects` by which backtracking is allowed. Planning here is searching through a dynamically constructed AND/OR graph. and *Provisional Agreement Protocol* (PAP) [Perugini *et al.*, 2004] also allows `provisional bid` and `withdraw bid`. ECNP has several applications in the military logistics. ECNP has been connected with linear programing and trust oriented reasoning [Rehak *et al.*, 2006]. Neither ECNP/PAP allow for work in semi-trusted environment with partially undisclosed information.

There are various approaches to task allocation in the open, semi-trusted communities. Unlike in the cooperative domains, in semi-trusted domains, the exact information about agents status, its' resource availability, set of agents' commitments and its strategies are not available. An example of such a method is an efficient algorithm for approximating the agents' equilibrium strategies [Sarne and Kraus, 2005]. In [Pechoucek *et al.*, 2006] there is presented a method for incremental negotiation based on approximation of the agents resource availability. The approximation is improved by means of the information included in unsuccessful negotiation rounds.

2.3.3 Enforcement of Coordinated Action

As already noted, It is important to make sure that the distributed plans will be robust and stable even in the self-interested communities and very dynamic environments. Multi-agent community provides mechanisms for specifying formal models of agents *commitments* as mental structures in their programmes [Excelente-Toledo *et al.*, 2001; Philippe Pasquier, 2004]. An inseparable part of each commitment is a specification of the conditions/postconditions under which the agents are allowed to drop their commitments. There is different use of commitments in the collaborative and competitive environments. While in cooperative settings the commitments postconditions provide mainly notification functionality, in the competitive environment the commitment postcondition provide incentive for an agent to keep its commitment (mainly in the from of penalties). It is believed that the combination of both would be necessary in the presented scenario, mainly for implementation of the replanning functionality.

2.3.4 Plan Execution Monitoring and Repair

AI planning systems take as input an abstract description of a planning domain. Since an abstraction is almost by definition an inaccurate model, it has to be

[3] i.e. they provide a solution that is independent of the order of the negotiation round

expected that long plans may fail at some point during the execution.

There are a number of reactive plan execution frameworks that include a monitoring component that traces the status of the execution to determine when the execution of an action does not lead to the expected result or assumptions made at plan generation time may not hold [Muscettola *et al.*, 1998; Myers and Wilkins, 1998]. When the observed state does not match expectations more planning is initiated.

There are two basic ways in which the planner can come up with a new plan that leads from the current situation (after some failure has occurred) to the desired goals. Firstly, the planner can abandon the current plan completely a generate a new plan from scratch which is known as replanning, or the planner can try to modify the remainder of the failed plan in various ways such that the modified plan is applicable and achieves the goals. The traditional way to decide what is is best is to ask which approach is more efficient. This question has been addressed from a theoretical perspective [Nebel and Koehler, 1995] as well as in empirical studies [Gerevini and Serina, 2000]. While efficiency is an important concern, it is not the only one. For example, [Fox *et al.*, 2006] have recently looked at plan stability which is important in many contexts, e.g. when the tasking of agents execution the plan cannot be changed easily or only at a cost. To this end they have modified an existing planner by incorporating plan stability into the evaluation function for a local search. The results show that similar plans can be found as the result of a repair, and that the performance is still good. How important plan stability is must depend on the domain though.

3 Integrated Architecture

I-Globe, the integrated technology of **a/globe** multi-agent technology and I-X task-support and planning technology, intends to support the process of distributed planning and coordination of team-oriented activities at three levels. Prior to discussing these levels, let us detail the technologies that the partners are bringing into the project and outline the scenario in which they will be applied.

AIAI is bringing to the project artificial intelligence based technology to provide task support

- I-X technology including the I-Plan planning agent which is based on a combination of a human-relatable Hierarchal Task Network (HTN) approach coupled with rich constraint representation and satisfaction algorithms.
- `<I-N-C-A>`(Issues, Nodes, Constraints and Annotations) as a shared ontology suitable for relating the activities of human, vehicle, robots and sensors.

- Previous relevant work to connect I-X/I-Plan to strategic issue-based decision support tools, such as Compendium, which have been demonstrated in DARPA programs such as Integrated Battle Command.

ATG is providing the following pieces of agent-based technology:

- **a/globe** multi-agent integration platform: This technology provides means for linking heterogenous software applications and reasoning algorithms. In addition to this, **a/globe** facilitates modeling of the environment in which the autonomous actors interact as well as modeling of the behavior of the actors themselves. Therefore **a/globe** is an ideal technology for development of the proof-of-concept prototype of the targeted scenario. **a/globe** also includes basic tools for 2D and 3D visualization support.

- **stand-in agent**: We provide technology supporting agents' interaction in the situation of temporal communication inaccessibility caused by either unreliable communication infrastructure, dynamics of an ad-hoc networking environment or actors altering offline/online status.

- **agent based planning**: ATG provides tactical planning and resource allocation algorithm based on ECNP (extended-contract-net-protocol) negotiation algorithm for distributed environment, that has been successfully deployed in military logistics scenarios.

- if the final demo requires, ATG provides other agent-based algorithms for handling and sharing semi-private information, representation and modeling of trust and reputation and techniques for creating social models of the individual actors.

3.1 I-Plan

The facilities available in the I-X Process Panels include an AI planner (I-Plan) used to provide context sensitive options for the handing of issues (such as the achievements of stated objectives), the performance of activities, and the satisfaction of constraints i.e. to support the underlying `<I-N-C-A>` plan representation.

For any activity on the panel, an "Action" column shows its current status and the available options to perform the activity. Colours indicate the readiness of the item for current execution.

- White indicates that the item is not currently ready for execution (i.e., some temporal ordering, preconditions or other constraints might not be met).

- Orange indicates that the action is ready to perform and that all preconditions and constraints are met.

- Green indicates that the item is currently being performed.

– Blue indicates successful completion.

– Red indicates a failure for which failure recovery planning steps might be initiated.

The set of "Actions" available to perform any item on the panel is available through a menu. This is dynamically generated and context-sensitive reflecting the knowledge of the capabilities of other panels and services available. It also draws on the in-built planner I-Plan to select from any known plans or Standard Operating Procedures ("plan schemas") that match the item.

I-Plan can perform hierarchical partial-order composition of plans from a library of single level plan schemas or "Standard Operating Procedures". This library can be augmented during planning either with a simple activity details interface to add in specific ways to expand a given action (intended for use by users familiar with the application domain but not AI planning techniques) or with a more comprehensive graphical domain editor. Grammars and lexicons for the domain are built automatically during domain editing to assist the user.

Future developments of I-Plan will be able to account for plan repair after partial failures, include handlign of calendar time and consumable resource contraints, and account for mutual satisfaction of open variables and other constraints with greater efficiency.

3.2 Agent-based Planning

In ATG we have developed a number different **a/globe** integrated planning algorithms and other auxiliary mechanisms that support the distributed planning scenario.

a/globe provides negotiation and auctioning mechanisms that facilitate distributed resource allocation. In **a/globe** there is also integrated the ECNP auctioning mechanism that provides optimization of multiple rounds in the negotiation process. Performance of ECNP has been already verified on the ground logistics scenario [Rehak *et al.*, 2006]. Therefore the requirements for decentralization [R1] of the planning process can be covered by **a/globe**.

a/globe also hosts agents with rich social models, the containers of agents mutual awareness. The social models are built from analysis of agents past interaction and their previous behavior. The social models can be used in the situations with partial knowledge availability or temporal communication inaccessibility (as specified by the requirements [R2] and [R3]. The current implementation of the social models is not very expressive and would need to be extended. Preliminary experiences verified the use of social model in semi-trusted logistics environment and the methods of incremental construction of the social models were deployed [Pechoucek *et al.*, 2006]. However we need to build methods that would allow the agents to share their information about the quality of the social models and would alow their run-time integration.

Similarly the methods of plan merging and agents' peer-to-peer negotiation about incomplete partial plans has not been yet developed in **a/globe**.

Key research and technological challenges are in linking the classical planning functionalities of e.g. HTN planning with the very dynamic nature of the multi-agent simulation. This needs to be done for support of not only the [R1] requirement but mainly for the [R4] requirement for dynamics and [R5] requirement for opportunistic planning.

a/globe integrates the classical A^*-based path planning mechanisms that are used for the **a/globe**-based application for UAV collision avoidance. While these algorithms do not support even the tactical planning level (discussed in Section 1.2), they will be critical for some aspects of the final demonstration development.

4 List of Deployment Scenarios

A simple emergency response domain model and set of tasks have been used for the initial work to connect I-Plan and **a/globe** to prove feasibility of the approach. This domain was specifically designed to avoid a lot of work in changing aspects of I-Plan and **a/globe** at this initial stage, so that existing example domains working already in each could be used. The specific scenarios that were worked on in the project were examples that seemed to fit well to allow us to test the link between I-Plan and **a/globe**. But of course these are quite simple.

4.1 Phase 1: Bus Crash Scenario

The simplest scenario has been reported in the interim report and is implemented on top of **a/globe** multi-agent infrastructure and existing implementation of ACROSS scenario, both provided by ATG. AIAI is has provided the existing HTN planner – I-Plan and a set of I-X Panels for human-machine interaction.

The first phase integration scenario, demonstrating full interoperability between **a/globe** and I-X has been decided upon and it goes as follows:

Let us have ACROSS scenario. The simulator will create additional "injured person" events that get sent to the agent containing the I-X panel and planner. The sole parameter is the location of the injured person. When the I-X agent receives this it will generate a plan consisting of three activities:

1. transport a doctor to the location of the location of the injury
2. transport some medical equipment to the location of the injury
3. transport the injured person to a hospital

The first two activities are not ordered, but the third activity has to occur after the first two. The plan will not contain the starting points for the trucks transporting the doctor of the medical equipment, and neither will it specify the hospital which is the destination of the third step. This plan will be sent to another **a/globe** agent that tries to allocate the three activities using the CNP just as it would for grain. Initially this will succeed and there will be no need for re-planning.

Figure 1: Common traffic with a bus crash and a terrorist (Terrorist2) leaving the accident on the right side, and on the left a police team patrolling. We can also see the cities and hospitals.

The point of this integration is to show that messages generated from **a/globe** can be translated into <I-N-C-A>-style messages that can be handled by an I-X agent, and that the plan generated by the I-X agent can be sent to an **a/globe** agent which negotiates and assigns other agents to the activities in the plan (initially only transportation activities) and manages the plan execution.

The aim of the phase 1 scenario was to show that the two frameworks, **a/globe** and I-X could be integrated and perform the appropriate message exchanges that were necessary for the later stages. As such, it was intended to provide a very simple test scenario for the infrastructure that had to be put in place. The served to highlight a number of issues that were subsequently resolved to achieve the first integration.

This scenario has been successfully developed during the first phase of the project – see the interim report.

4.2 Phase 2: Extended Scenario

After successful implementation of the first phase, aim of which was to prove the compatibility of the two systems, we have concentrated on taking advantage of the integrated system and implemented an extended scenario. In this scenario we are demonstrating these aspects of the system:

– **Distributed planning, scheduling and execution of plans**, which controls the rescue actions. There are more emergency centers in the simulation and they share the resources. There are also more planners in the system that can produce different plans for the same accident. The scheduling is distributed as well and uses the Contract-Net Protocol.

– **Extended planner domain**, which can produce various plans for handling the problems so that the rescue action does not fail so easily e.g. after one resource could not be scheduled.

– **Extended simulation abilities**, which include simulation of various adversary behavior, different types of accidents/disasters, which have impact on the common traffic on the roads as well as on the cities, different types of accidents' handling – the rescue units can have variety of abilities e.g. rescuing wounded, fighting fire, quelling insurgency etc.

4.2.1 Entities in the Simulation

Figure 2: Two disasters created by the terrorists. There are wounded people and also there is a potential for spreading the insurgency. The common traffic is not shown in the picture for the sake of clarity.

Let us first describe the entities that take part in the simulation:

– *Emergency centers* – agents responsible for handling the emergencies like e.g. bus crashes, insurgencies and fires. The centers are collecting the calls from the other entities and creating rescue actions.

– *Emergency units* – ambulances, police teams, trucks with material. These entities carry out the rescue actions. They can transport material and staff (doctors, paramedics) to and from the sites of accidents.

– *Adversaries* – entities that are attacking the soft targets like buses, setting cities on fire and destroying roads.

- *Hospitals, Depots* – places where the material and staff for handling the accidents is located, or where the wounded people are transported to.

- *Cities* – cities in the simulation are representing the civilians who demand goods. The traffic on the roads is actually the result of this demand, since the trucks are transporting goods between the cities.

- *Common traffic* – trucks that transport goods between cities. These entities are subject to adversary activity since the adversaries can block the roads and the trucks have to take longer routes to the delivery destinations.

- *Transporters* – the owners of the trucks. They make contracts with the cities for transport of goods.

We can see pretty much all the entities at Figure 1.

4.2.2 Scenario Description

Figure 3: The accident labeled BusCrash0 has been reported and the EmergencyCentre1 has become responsible for its handling.

At the beginning of the scenario the world is in state of peace. There is the common traffic on the roads going between the cities and transporting goods. There are also police teams cruising the roads and looking for possible accidents. However, there are also insurgents/terrorist that roam around the world and once in a while crash buses or set something at fire.

Once there is an accident the terrorists run away from it. They can be caught by the police teams if they meet, but there always stays at least one terrorist in the simulation world. We can see accidents' creation at Figure 2. The accidents force the common traffic to take another road to their destinations making them possibly late for their delivery times. If the accidents have certain insurgency level, they can also have a side effect – when left unattended for some time they can create another terrorist cell/entity in a nearby city.

Once any entity that is not being a terrorist or insurgent comes across an accident it reports it to at least one emergency center. If the accident is reported to more than one emergency center, they negotiate about who will handle the accident. The situation, right after an accident has been reported, can be seen at Figure 3.

After the emergency centers settle on who is going to handle the accident, the chosen center starts the rescue action. There are these phases of the rescue action (every phase has a color in the visualization that is displayed on the emergency center's chart next to the names of handled accidents):

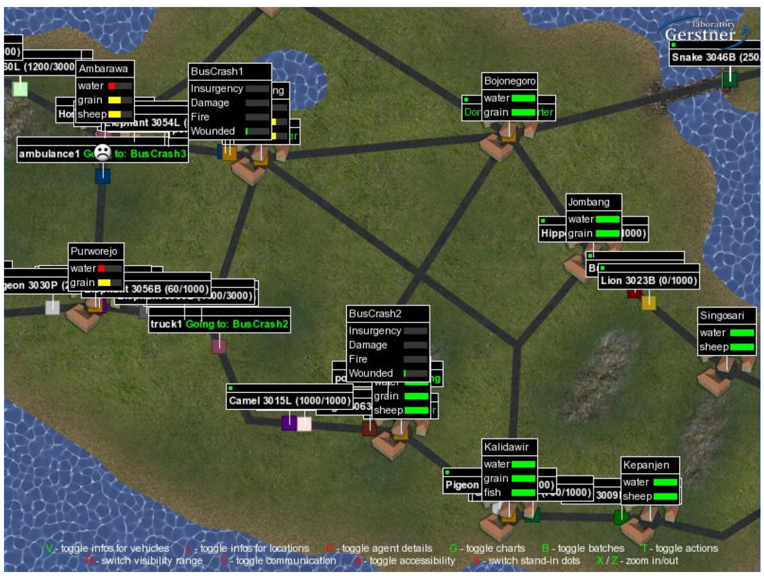

Figure 4: In the middle left there is a truck (truck1) going to an accident (BusCrash2) carrying doctors to take care of the wounded. This is the case of replanning, since the truck is used instead of an ambulance that would be preferred if available.

1. Planning (blue)– the emergency center at first has only the description of the accident. The agent representing the center passes the description to its I-X planner to get a plan of activities how to handle the accident. We can see a plan for an accident at Figure 5.

2. Scheduling (yellow) – once the emergency center has the plan it can start the scheduling procedure. Scheduling is done using the Contract-Net protocol and it tries to schedule the actors (police-teams, ambulances etc.) and the material (medical material) needed. When a resource is unavailable the plan has to be **replanned** (see Figure 4 for a case of replanning),

i.e. we pass the type of unavailable resource to the planner and we go back to phase one.

3. Execution (orange) – after successful scheduling of resources the execution starts. Every actor is given tasks to do and is responsible for reporting status of the tasks back to the emergency center for execution monitoring.

4. Done (green)/Failed (red) - if the accident is handled successfully it is marked done, logged and removed from active tasks of the emergency center. If the handling fails, the emergency center starts the whole procedure over.

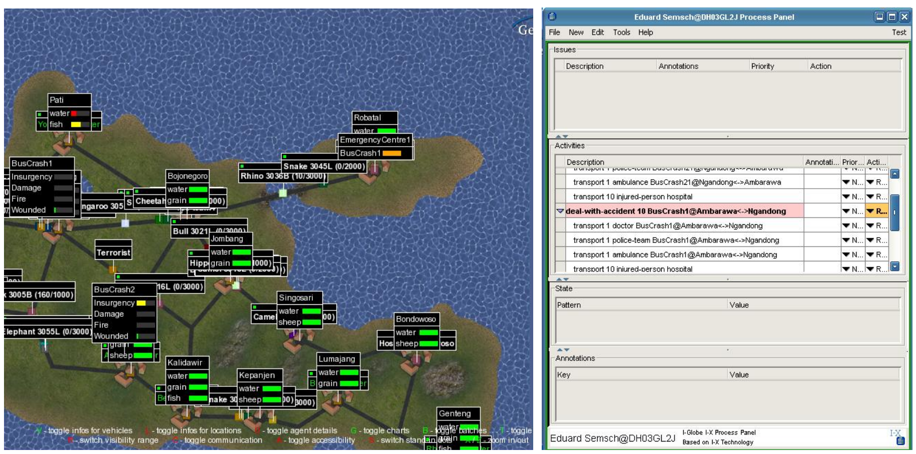

Figure 5: We see a bus crash on the left side and the respective plan for its handling in the planner window on the right.

4.3 Phase 3: Go Places and Do Things Application

However, for continuing work needed a good flexible open ended and simply structured domain into which we can introduce a wide range of incidents and a wide and interesting set of incident response resources and transports. We want to introduce multiple mobile vehicles and UAVs as well as more interesting concurrent incidents with contentions for resources. In preparation for this, we have adapted an O-Plan [Currie and Tate, 1991a] "Go Places Do Things" (GPDT) domain, which has the appropriate structure. We have translated this to be able to be used in I-Plan and initial experiments have been conducted between Edinburgh and Prague to connect the transportation planning level of this domain to the current **a/globe** flexible transportation and logistics planning facilities.

The next stage of the project is intended to use the GPDT domain. The aim here is to have a well structured domain that allows for any number of interesting and relevant extensions, including:

– Location and network of places and roads (with changing usability)

– Types of incident

– Types of transport and mobile vehicle (trucks, buses, UAVs, etc.)

– Types of indecent response resource (medical teams, rescue personnel, specialized equipment, etc.)

There is an island with 5 cities (Abyss, Barnacle, Calypso, Delta, Exodus), and connecting roads. The cities have different populations ranging from 100 to 3000. Note the these are the cities in the present scenario that is set on a fictional island, called Pacifica.

There are different kinds of emergency response teams that can respond to different kinds of problems, e.g. medical teams can deal with injured people, evacuation teams can evacuate populations, etc. Some teams require equipment to do their work (e.g. medical equipment) and different kinds of supplies are available if required (e.g. food supplies).

When an emergency occurs, an appropriate and available team is selected and sent to the location of the problem to perform their response tasks. Any equipment and supplies that are required will also be sent to the location of the problem as part of the response.

Some teams can deal with different types of problems, e.g. a medical team can also organise an evacuation if it is dealing with injured at the location, but a medical team would not normally be chosen to evacuate a place when there are no injured people.

The response-agent implemented in this scenario has no transport layer. Transport is provided by a different agent with activities of the form `(transport ?number-of-items ?item-type ?to-location)`. Knowledge of roads or vehicles is therefore not required in the response-agent.

5 Conclusion

It has been demonstrated that the close integration of the classical hierarchical task network (HTN) oriented planning technologies can be easily integrated with the state-of-the-art agent technology. While HTN planning supports primarily the strategic and operational level, the multi-agent planning techniques play an important role in operational and tactical planning.

Agent technology also provides an efficient approach to simulation of plan execution. Thus we see a great potential of the presented integrated architecture especially in the tasks that related to real-time replanning in dynamic environment.

A Running the Demo

There are two programs needed for the demo. The first one is the simulation and the other one is the visualization. There are two configurations of the simulation - one (run the `action.bat` from simulation main folder) has a lot of entities and there are rescue actions, adversaries and common traffic, the other (run `straight.bat`) has only the rescue actions and adversaries to make the simulation more clear. It takes these steps to run the Demo:

1. Run the visualization first using the `across2d.exe` (preferably at the same computer, otherwise you need to change the configuration in `xml/3dvisio.xml`.

2. Run either of the `.bat` files (it is easy to change them to run in linux).

3. You should get five Java windows on the screen. Click the `EntMan` (entity manager) and press Start Scenario at the bottom of the window.

4. Once the I-X Process Panels (there are usually more planners in the scenario) are up, set them to the automatic mode (Using `Test` at the right upper corner) so that you don't have to plan and send the plan after every accident.

5. You may want to zoom the visualization (with X and Y on keyboard) or toggle communication (C) or toggle charts (G) so you can see it more clearly. Or you can pause the simulation using `EntitySimulator` window.

Running only the I-X part of the Demo

To start the demo in Windows: double-click on `apps/gpdt/scripts/win/1-response.bat`
To start the demo in Unix: ensure you have a `ix.jar` in your I-X base directory (the one that the apps directory is in). Also ensure that the unix script is executable (`chmod ugo+x apps/gpdt/scripts/unix/1-response`). Cd to the apps/gpdt directory, then from the unix prompt type `./scripts/unix/1-response`

You should now have an Emergency Response Panel which some state information and a response activity. Refine any orange activities on the list, binding variables as requested (when the activity turns pink, (right-click on the activity and select bindings)

Try complicating the emergency by adding an evacuation problem some of the way through (select add-problem from the Test menu, then choose an evacuation refinement) - the medical team should be able to cover the evacuation, but only after the injured are dealt with and if they are still assigned to the emergency when the evacuation problem occurs.

References

[Albore and Bertoli, 2006] Alexandre Albore and Piergiorgio Bertoli. Safe LTL assumption-based planning. In Derek Long, Stephen F. Smith, Daniel Borrajo, and Lee McCluskey, editors, *Proc. 16th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 193–202. AAAI Press, June 2006.

[Allen *et al.*, 1990] James Allen, James Hendler, and Austin Tate, editors. *Readings in Planning*. Morgan Kaufman, 1990.

[Bonet, 2006] Blai Bonet. Bounded branching and modalities in non-deterministic planning. In Derek Long, Stephen F. Smith, Daniel Borrajo, and Lee McCluskey, editors, *Proc. 16th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 42–52. AAAI Press, June 2006.

[Boutilier *et al.*, 1999] Craig Boutilier, Moises Goldszmidt, and Bikash Sabata. Sequential auctions for the allocation of resources with complementarities. In *IJCAI '99: Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pages 527–523. Morgan Kaufmann Publishers Inc., 1999.

[Brafman and Chernyavsky, 2005] R. Brafman and Y. Chernyavsky. Planning with goal preferences and constraints. In *Proc. 15th International Conference on Automated Planning and Scheduling (ICAPS)*, 2005.

[Bryce *et al.*, 2006] Daniel Bryce, Subbarao Kambhampati, and David Smith. Sequential Monte Carlo in probabilistic planning reachability heuristics. In Derek Long, Stephen F. Smith, Daniel Borrajo, and Lee McCluskey, editors, *Proc. 16th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 233–242. AAAI Press, June 2006.

[Castillo *et al.*, 2006] Luis Castillo, Juan Fdez-Olivares, Oscar Garcia-Perez, and Francisco Palao. Efficiently handling temporal knowledge in an HTN planner. In Derek Long, Stephen F. Smith, Daniel Borrajo, and Lee McCluskey, editors, *Proc. 16th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 63–72. AAAI Press, June 2006.

[Chang and Amir, 2006] Allen Chang and Eyal Amir. Goal achievement in partially known, partially observable domains. In Derek Long, Stephen F. Smith, Daniel Borrajo, and Lee McCluskey, editors, *Proc. 16th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 203–211. AAAI Press, June 2006.

[Chen *et al.*, 1999] Y. Chen, Y. Peng, T. Finin, Y. Labrou, and S. Cost. A negotiation-based multi-agent system for suppy-chain management. In *Proc. Workshop on Agent-Based Decision Support in Managing the Internet-Enabled Supply-Chain at Agents'99*, 1999.

[Cox and Durfee, 2005] Jeffrey S. Cox and Edmund H. Durfee. An efficient algorithm for multiagent plan coordination. In *AAMAS '05: Proceedings of the fourth international joint conference on Autonomous agents and multiagent systems*, pages 828–835, New York, NY, USA, 2005. ACM Press.

[Currie and Tate, 1991a] K. Currie and A. Tate. O-plan: the open planning architecture. *Artificial Intelligence*, 52:49–86, 1991.

[Currie and Tate, 1991b] Ken Currie and Austin Tate. O-Plan: The open planning architeture. *Artificial Intelligence*, 52:49–86, 1991.

[Dadeh *et al.*, 1999] N. Dadeh, D. Hildum, D. Kjenstad, and A. Tseng. Mascot: An agent-based architecture for coordinated mixed-initiative supply chain planning and scheduling. In *Proc. Workshop on Agent-Based Decision Support in Managing the Internet-Enabled Supply-Chain at Agents'99*, 1999.

[Dechter, 2003] R. Dechter. *Constraint Processing*. Morgan Kaufman, 2003.

[Doucet *et al.*, 2001] A. Doucet, N. de Freita, and N. Gordon. *Sequential Monte Carlo Methods in Practice*. Springer, 2001.

[Drabble and Tate, 1994] Brian Drabble and Austin Tate. The use of optimistic and pessimistic resource profiles to inform search in an activity based planner. In *Artificial Intelligence Planning Systems*, pages 243–248, 1994.

[El Sakkout and Wallace, 2000] Hani El Sakkout and Mark Wallace. Probe backtrack search for minimal perturbation in dynamic scheduling. *Constraints*, 5(4):359–388, 2000.

[Even-Dar *et al.*, 2005] E. Even-Dar, S. Kakade, and Y. Mansour. Reinforcement learning in POMDSs. In *Proc. 19th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 660–665. AAAI Press, 2005.

[Excelente-Toledo *et al.*, 2001] C.B. Excelente-Toledo, R.A. Bourne, and N.R. Jennings. Reasoning about commitments and penalties for coordination between autonomous agents. In *Proceedings of the Fifth International Conference on Autonomous Agents, 28 May-1 June 2001, Montreal, Que., Canada*, pages 131–8, New York, USA,, 2001. ACM.

[Fdez-Olivares *et al.*, 2006] Juan Fdez-Olivares, Luis Castillo, Oscar Garcia-Perez, and Francisco Palao. Bringing users and planning technology together: Experiences in SIADEX. In Derek Long, Stephen F. Smith, Daniel Borrajo, and Lee McCluskey, editors, *Proc. 16th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 11–20. AAAI Press, June 2006.

[Fischer *et al.*, 1995] Klaus Fischer, Jörg P. Muller, Markus Pischel, and Darius Schier. A model for cooperative transportation scheduling. In *Proceedings of the First International Conference on Multiagent Systems.*, pages 109–116, Menlo park, California, June 1995. AAAI/MIT Press.

[Fox and Long, 2002] Maria Fox and Derek Long. PDDL2.1: An extension to pddl for expressing temporal planning domains. http://www.dur.ac.uk/d.p.long/competition.html, 2002.

[Fox *et al.*, 2006] Maria Fox, Alfonso Gerevini, Derek Long, and Ivan Serina. Plan stability: Replanning versus plan repair. In Derek Long, Stephen F. Smith, Daniel Borrajo, and Lee McCluskey, editors, *Proc. 16th International*

Conference on Automated Planning and Scheduling (ICAPS), pages 212–221. AAAI Press, June 2006.

[Frank *et al.*, 2006] Jeremy Frank, James Crawford, Lina Khatib, and Ronen Brafman. Tractable optimal competitive scheduling. In Derek Long, Stephen F. Smith, Daniel Borrajo, and Lee McCluskey, editors, *Proc. 16th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 73–82. AAAI Press, June 2006.

[Fung and Chen, 2005] R.Y.K. Fung and T. Chen. A multiagent supply chain planning and coordination architecture. *International Journal of Advanced Manufacturing Technology*, 25(7-8):811–819, 2005.

[Gallager *et al.*, 2006] Anthony Gallager, Terry Zimmerman, and Stephen Smith. Incremental scheduling to maximize quality in a dynamic environment. In Derek Long, Stephen F. Smith, Daniel Borrajo, and Lee McCluskey, editors, *Proc. 16th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 222–231. AAAI Press, June 2006.

[Georgeff, 1987] Michael Georgeff. Planning. *Ann. Rev. Comput. Sci.*, 2:359–400, 1987. reprinted in [Allen *et al.*, 1990, pages 5–25].

[Gerevini and Serina, 2000] Alfonso Gerevini and Ivan Serina. Fast plan adaptation through planning graphs: Local and systematic search techniques. In *Proc. International Conference on AI Planning and Scheduling (AIPS)*, pages 112–121, 2000.

[Ghallab *et al.*, 2004] Malik Ghallab, Dana Nau, and Paolo Traverso. *Automated Planning.* Morgan Kaufman, 2004.

[Green, 1969] Cordell Green. Application of theorem proving to problem solving. In *Proc. 1st International Joint Conference on Artificial Intelligence (IJCAI)*, pages 741–747, 1969. reprinted in [Allen *et al.*, 1990, pages 67–87].

[Hunsberger and Grosz, 2000] L. Hunsberger and B. J. Grosz. A combinatorial auction for collaborative planning. In *ICMAS '00: Proceedings of the Fourth International Conference on MultiAgent Systems (ICMAS-2000)*, pages 151–158. IEEE Computer Society, 2000.

[Hunsberger, 2002] Luke Hunsberger. *Group Decision Making and Temporal Reasoning.* PhD thesis, Harvard University, 2002.

[Kearns *et al.*, 2000] M. Kearns, Y. Mansour, and A.Y. Ng. Approximate planning in large POMDPs via reusable trajectories. In *Proc. 19th Conference on Neural Information Processing Systems (NIPS)*, pages 1001–1007. MIT Press, 2000.

[Lawton and Domshlak, 2004] James H. Lawton and Carmel Domshlak. Multiagent opportunistic planning and plan execution. In *ICTAI '04: Proceedings* *of the 16th IEEE International Conference on Tools with Artificial Intelligence (ICTAI'04)*, pages 408–415, Washington, DC, USA, 2004. IEEE Computer Society.

[McDermott, 1998] Drew McDermott. PDDL: The planning domain definition language. Technical report, Yale Center for Computational Vision and Control, 1998. ftp://ftp.cs.yale.edu/pub/mcdermott/software/pddl.tar.gz.

[Meuleau *et al.*, 2006] Nicolas Meuleau, Ronen Brafman, and Emmanuel Benazera. Stochastic over-subscription planning using hierarchical MDPs. In Derek Long, Stephen F. Smith, Daniel Borrajo, and Lee McCluskey, editors, *Proc. 16th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 121–130. AAAI Press, June 2006.

[Morris *et al.*, 2004] P. Morris, R. Morris, L. Khatib, S. Ramakrishnan, and A. Bachmann. Strategies for global optimization of temoral preferences. In *Proc. 10th International Conference on Principles and Practices of Constraint Programming*, pages 408–422, 2004.

[Muscettola *et al.*, 1998] N. Muscettola, P. Nayak, B. Bell, and B. Williams. Remote agent: To boldly go where no AI system has gone before. *Artificial Intelligence*, 103:5–47, 1998.

[Myers and Wilkins, 1998] Karen Myers and David Wilkins. Reasoning about locations in theory and practice. *Computational Intelligence*, 14:151–187, 1998.

[Nau *et al.*, 2001] D. Nau, H. Munoz-Avila, Y. Cao, A. Lotem, and S. Mitchell. Total-order planning with partially ordered subtasks. In *Proc. 17th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 425–430. AAAI Press, 2001.

[Nebel and Koehler, 1995] B. Nebel and J. Koehler. Plan reuse versus plan generation: A theoretical and empirical analysis. *Artificial Intelligence*, 76(1):427–454, 1995.

[Pardoe and Stone, 2006] David Pardoe and Peter Stone. Predictive planning for supply chain management. In Derek Long, Stephen F. Smith, Daniel Borrajo, and Lee McCluskey, editors, *Proc. 16th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 21–30. AAAI Press, June 2006.

[Pěchouček *et al.*, 2002] Michal Pěchouček, Vladimír Mařík, and Jaroslav Bárta. A knowledge-based approach to coalition formation. *IEEE Intelligent Systems*, 17(3):17–25, 2002.

[Pechoucek *et al.*, 2006] Michal Pechoucek, Ondrej Lerch, and Jiri Biba. Iterative query-based approach to efficient task decomposition and resource allocation. In Matthias Klusch, Michael Rovatsos, and Terry Payne, editors, *Cooperative Information Agents X*, volume 4149 of *Lecture Notes in Computer Science*, pages 258–272. CIA Workshop, Springer, 2006.

[Pednault, 1989] E. Pednault. ADL: Exploring the middle ground between STRIPS and the situation calculus. In *Proc. International Conference on Knowledge Representation and Reasoning (KR)*, pages 324–332, 1989.

[Pednault, 1994] E. Pednault. ADL and the state-transition model of action. *Journal of Logic and Computation*, 4:467–512, 1994.

[Perugini *et al.*, 2004] Don Perugini, Dale Lambert, Leon Sterling, and Adrian Pearce. Provisional agreement protocol for global transportation scheduling. In *Workshop on agents in traffic and transportation held in conjunction with the International Conference on Autonomous Agents and Multi Agent Systems (AAMAS04)*, New York, 2004.

[Philippe Pasquier, 2004] Brahim Chaib-draa Philippe Pasquier, Roberto Flores. The enforcement of flexible social commitments. In F. Zambonelli M.-P. Gleizes, A. Omicini, editor, *Fifth International Workshop Engineering Societies in the Agents World (ESAW), Lecture Notes in Artificial Intelligence (LNAI),*, volume 3451, pages 139–152. Springer-Verlag, 2004.

[Reece and Tate, 1994] Glen A. Reece and Austin Tate. Synthesizing protection monitors from causal structure. In *Artificial Intelligence Planning Systems*, pages 146–151, 1994.

[Rehak *et al.*, 2006] Martin Rehak, Michal Pechoucek, and Premysl Volf. Multilevel approach to agent-based task allocation in transportation. In Matthias Klusch, Michael Rovatsos, and Terry Payne, editors, *Cooperative Information Agents X*, volume 4149 of *Lecture Notes in Computer Science*, pages 273–287. Springer, 2006.

[Sanchez and Kambhampati, 2005] R. Sanchez and S. Kambhampati. Planning graph heuristics for selecting objectives in over-subscription planning problems. In *Proc. 15th International Conference on Automated Planning and Scheduling (ICAPS)*, 2005.

[Sandholm and Suri, 2000] T. Sandholm and S. Suri. Improved algorithms for optimal winner determination in combinatorial auctions and generalizations. In *Proc. 17th National Conference on Artificial Intelligence*, pages 90–96. AAAI Press, 2000.

[Sandholm, 2002] T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135(1-2):1–54, 2002.

[Sarne and Kraus, 2005] David Sarne and Sarit Kraus. Solving the auction-based task allocation problem in an open environment. In *Proceedings of Twentieth National Conference on Artificial Intelligence*, pages 164–169, 2005.

[Siebra and Tate, 2006] Clauirton Siebra and Austin Tate. An investigation into the use of collaborative concepts for planning in disaster response coalitions. In *DIS '06: Proceedings of the IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and Its Applications (DIS'06)*, pages 253–258, Washington, DC, USA, 2006. IEEE Computer Society.

[Smith and Weld, 1998] D. Smith and D. Weld. Conformant Graphplan. In J. Mostow and C. Rich, editors, *Proc. 15th National Conference on Artificial Intelligence (AAAI)*, pages 889–896. MIT Press/AAAI Press, 1998.

[Smith, 1980] R. G. Smith. The contract net protocol: High level communication and control in a distributed problem solver. *In IEEE Transactions on Computers*, C-29(12):1104–1113, 1980.

[Smith, 1994] Stephen Smith. OPIS: A methodology and architecture for reactive scheduling. In M. Zweben and M. Fox, editors, *Intelligent Scheduling*. Morgan Kaufmann, 1994.

[Smith, 2004] D. Smith. Choosing objectives in over-subscription planning. In *Proc. 14th International Conference on Automated Planning and Scheduling (ICAPS)*, pages 393–401, 2004.

[Tate and Selvin, 2006] Buckingham Shum S.J. Dalton J. Mancini C. Tate, A. and A.M. Selvin. Co-opr: Design and evaluation of collaborative sensemaking and planning tools for personnel recovery, 2006.

[Tate and Stader, 2002] Dalton J. Tate, A. and J. Stader. I-p2 - intelligent process panels to support coalition operations. In *Proceedings of the Second International Conference on Knowledge Systems for Coalition Operations*, 2002.

[Tate *et al.*, 1994] A. Tate, B. Drabble, and J. Dalton. Reasoning with constraints within o-plan, 1994.

[Tate *et al.*, 1999] A. Tate, J. Levine, J. Dalton, and S. Aitken. Supporting the planning process using open planning process panels, 1999.

[Tate, 1984] Austin Tate. Planning and condition monitoring in a fms. 1984.

[Tate, 1989] A. Tate. Coordinating the activities of a planner and an execution agent. In G. Rodriguez, editor, *Proc. NASA conf. on space telerobotics*, NASA JPL, 1989. JPL Publications.

[Tate, 1993] A. Tate. Putting knowledge-rich plan representations to use, 1993.

[Tate, 1996] Austin Tate. Using constraints for task-oriented communication, planning and control. 1996.

[Tate, 2000a] A. Tate. <i-n-ova>and <i-n-c-a> - representing plans and other synthesized artifacts as a set of constraints. In *AAAI-2000 Workshop on Representational Issues for Real-World Planning Systems*, 2000.

[Tate, 2000b] A. Tate. Intelligible ai planning, in research and development in intelligent systems. In *Proceedings of ES2000*. Springer, 2000.

[Tate, 2003] Austin Tate. I-x and <i-n-c-a>: an architecture and related ontology for mixed-initiative synthesis tasks, 2003.

[Tonino *et al.*, 2002] J.F.M. Tonino, André Bos, Mathijs M. de Weerdt, and Cees Witteveen. Plan coordination by revision in collective agent-based systems. *Artificial Intelligence*, 142(2):121–145, 2002.

[Valk *et al.*, 2005] Jeroen M. Valk, Mathijs M. de Weerdt, and Cees Witteveen. Algorithms for coordination in multi-agent planning. In Vlahavas and Vrakas, editors, *Intelligent Techniques for Planning*, pages 194–224. Idea Group Inc., 2005.

[van der Krogt *et al.*, 2003] R.P.J. van der Krogt, M.M. de Weerdt, and C. Witteveen. A resource based framework for planning and replanning. In *Web Intelligence and Agent Systems*, volume 1, pages 173–186. IOS Press, 2003.

[von Martial, 1992] Frank von Martial. *Coordinating Plans of Autonomous Agents*, volume 610 of *Lecture Notes on Artificial Intelligence*. Springer Verlag, Berlin, 1992.

[Šišlák *et al.*, 2006] David Šišlák, Martin Rehák, Michal Pěchouček, and Dušan Pavlíček. Deployment of a-globe multi-agent platform. In *AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems*, pages 1447–1448, New York, NY, USA, 2006. ACM Press.

[Wickler, 2006] Potter S. Tate A. Wickler, G. Using i-x process panels as intelligent to-do lists for agent coordination in emergency response. *International Journal of Intelligent Control and Systems (IJICS), Special Issue on Emergency Management Systems*, 2006.

[Wilkins, 1988] David Wilkins. *Practical Planning: Extending the Classical AI Planning Paradigm*. Morgan Kaufman, 1988.

[Witteveen and de Weerdt, 2006] Cees Witteveen and Mathijs de Weerdt. Multi-agent planning for non-cooperative agents. In Ed Durfee and David Musliner, editors, *Proceedings of the AAAI Spring Symposium on Distributed Plan and Schedule Management*, pages 169–170. AAAI, AAAI Press, 2006.

[Yang, 1997] Qien Yang. *Intelligent Planning*. Springer-Verlag, Berlin, Germany, 1997.