

Coordinating Planning Activity and Information Flow in a Distributed Planning System

Marie desJardins and Michael Wolverton

SRI International
 333 Ravenswood Ave.
 Menlo Park CA 94025
 {marie,mjw}@ai.sri.com

Abbreviated title: Coordinating a Distributed Planning System

Abstract

Distributed SIPE (DSIPE) is a distributed planning system that provides decision support to human planners in a collaborative planning environment. The key contributions of our research on DSIPE are (1) constraint-based, consistent local views of the global plan that give each planner a view of how other planners' subplans relate to their local planning decisions; (2) methods for automatically identifying and sharing potentially relevant information among distributed planning agents; and (3) techniques for merging subplans that leverage the shared subplan structure to generate a complete, final plan. DSIPE is a fully implemented system, and has been demonstrated to end users in the maritime (U.S. Navy and U.S. Marine Corps) planning community.

Keywords: Distributed HTN planning, information management

Distributed Planning

This paper describes Distributed SIPE (DSIPE), a distributed version of the SIPE-2¹ hierarchical task network (HTN) planning system (Wilkins 1988). DSIPE provides the user with semi-automated generation of crisis response options, in the presence of multiple, competing objectives and constraints, within a distributed computing environment that includes multiple planners collaboratively creating a plan to achieve a set of common objectives. DSIPE is most closely related to Corkill's (1979) distributed version of NOAH, a non-linear hierarchical planner from which SIPE-2 is conceptually descended. DSIPE extends the ideas in Distributed NOAH by focusing on efficient communication among planners and on creating a common partial view of subplans. All of the capabilities described in this paper are part of the current implemented system, except where noted. Our current application domain is maritime campaign planning, but the techniques are domain-independent.

¹SIPE-2: System for Interactive Planning and Execution.

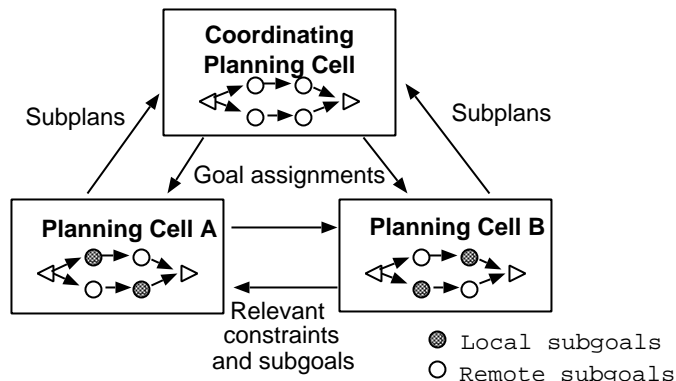


Figure 1: Distributed Planning Architecture

In our distributed planning environment, multiple copies of the DSIPE planner run on separate processors that are connected across a network. These distributed processes communicate with each other via message passing.² Each instance of DSIPE supports a human planner. We refer to the human planner as the *user*, to DSIPE as the *planning system*, and to the user and planning system together as a *planning cell*.

As in SIPE-2, plans in DSIPE are represented as partially ordered networks of tasks and subgoals. Plans are expanded by applying *planning operators*, which represent templates or strategies for solving a goal. Each operator has preconditions (gating applicability conditions), constraints (used for binding variables and assigning resources), a plot (a partially-ordered network of tasks and subgoals that expands the goal being solved), and effects (i.e., postconditions, which are expected to be true in the world state after the plot is applied).

A high-level view of our distributed planning architecture is shown in Figure 1. Goals are distributed

²Our system supports messages passed using the Knowledge Query Message Language (KQML) (Finin, Labrou, & Mayfield 1997), as text files, or as CORBA-compliant distributed objects, using Xerox PARC's ILU implementation (Janssen *et al.* 1998).

from a *coordinating planning cell* down to lower-level planning cells, which exchange information with each other as they expand their subplans. (Although only two lower-level planning cells are shown, the architecture can support an arbitrary number of cells. Also, a lower-level planning cell can itself act in the role of a coordinator for additional planning cells, resulting in a hierarchy of cells.) Each planning cell has a partial view of the other cells' subplans, with explicit dependencies and relationships to the local cell's subplan.

Throughout the paper, we demonstrate DSIPe's operation using an example from the maritime planning domain. The scenario is a noncombatant evacuation operation, in which civilian personnel must be evacuated from a city where there has been an insurrection. A Marine Corps planner and a Naval planner must work together to clear and land on a beach, move inland, and evacuate the civilians to an offshore vessel. The planners in this example have separate but related requirements, goals, and priorities. Although their requirements overlap for this mission, they each have independent goals (other missions to be performed or supported), capabilities, and resources.

Distributing Goals

Distributed planning in DSIPe follows the hierarchical model used by SIPE-2. The coordinating planning cell develops a high-level plan that achieves the top-level objectives of the current planning problem. The coordinator then distributes this high-level plan to lower-level planning cells, partitioning the subgoals among the cells. This model works well in the military environment, since the planning cell levels naturally correspond to echelons in the military hierarchy. In a less structured collaborative environment, alternative models of goal distribution might be required—for example, each cell might develop its own goal set and notify the other agents what it was planning to work on.

As shown in Figure 2, the planning cells expand their subplans separately, while sharing relevant information and constraints, as discussed in the following subsections. At the end of the planning process, the coordinator merges the subplans together, as discussed below in *Plan Merging*.

Figure 3 shows part of the Marine Corps subplan, before that planning cell has received any postconditions from the Naval planning cell. The hexagons represent goals, and the rounded oblongs represent primitive tasks. The arrows show temporal orderings, and activities that are not ordered may occur in parallel. In this subplan, the PREP-LAND (prepare to land on the beach) task is in parallel with another branch of the plan. In the latter (lower) branch, there are two goals in sequence, BEACH-APP-CLR (clear the approach to the beach, i.e., the water immediately adjacent to the beach, of mines and enemy troops) and RETURNED-TO (return to the offshore vessel). These two goals are in parallel with a third goal, BEACH-CLEARED (clear the beach itself). These two parallel

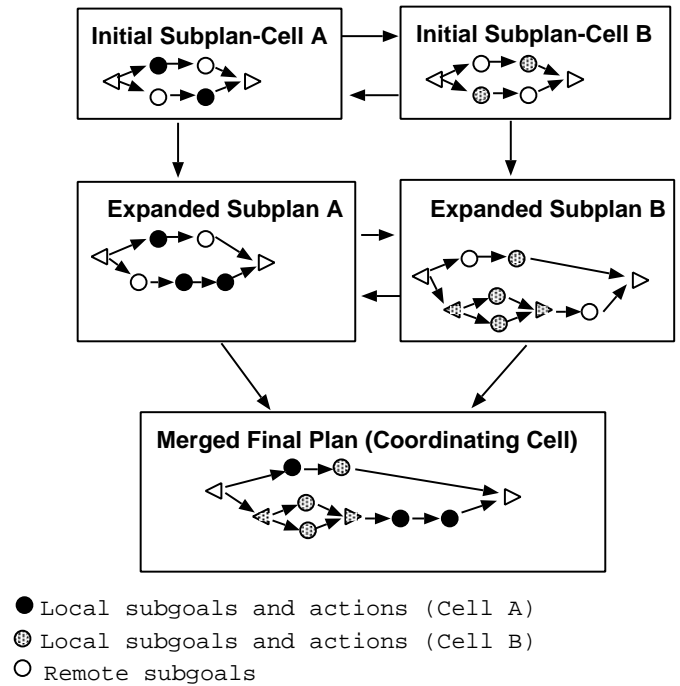


Figure 2: Expanding and Merging Subplans

branches are followed by several other goals, TRANS-TO (transport the Marines from the beach to the evacuation site), INTEL-DONE (complete an intelligence survey of the evacuation site), and EVAC-GRP-DONE (transport the evacuees to a safe location).

The darker nodes in the display correspond to goals in the remote (Naval) subplan. (These nodes appear in red in the screen display.) The other nodes are part of the local (Marine Corps) subplan. For example, the PREP-LAND task and the BEACH-CLEARED subgoal are part of the Marine Corps subplan; the BEACH-APP-CLR and RETURNED-TO subgoals belong to the Naval subplan.

Common Plan View

We have extended SIPE-2's internal plan representation so that each planning cell has a complete representation of its own subplan, as well as a partial representation of the subplans being developed by other planning cells. Subgoals being solved by other cells are represented in the local subplan as goal nodes that are not expanded. These nodes serve as placeholders for attaching postconditions and other constraints, which are sent by the remote planning cells that are responsible for expanding those goal nodes. The three types of constraints that are currently communicated among planning cells are *postconditions*, *ordering constraints*, and *variable constraints* (including temporal variables).

During plan generation, postconditions are associated with tasks and subgoals through the process of operator expansion. Each planning cell automatically

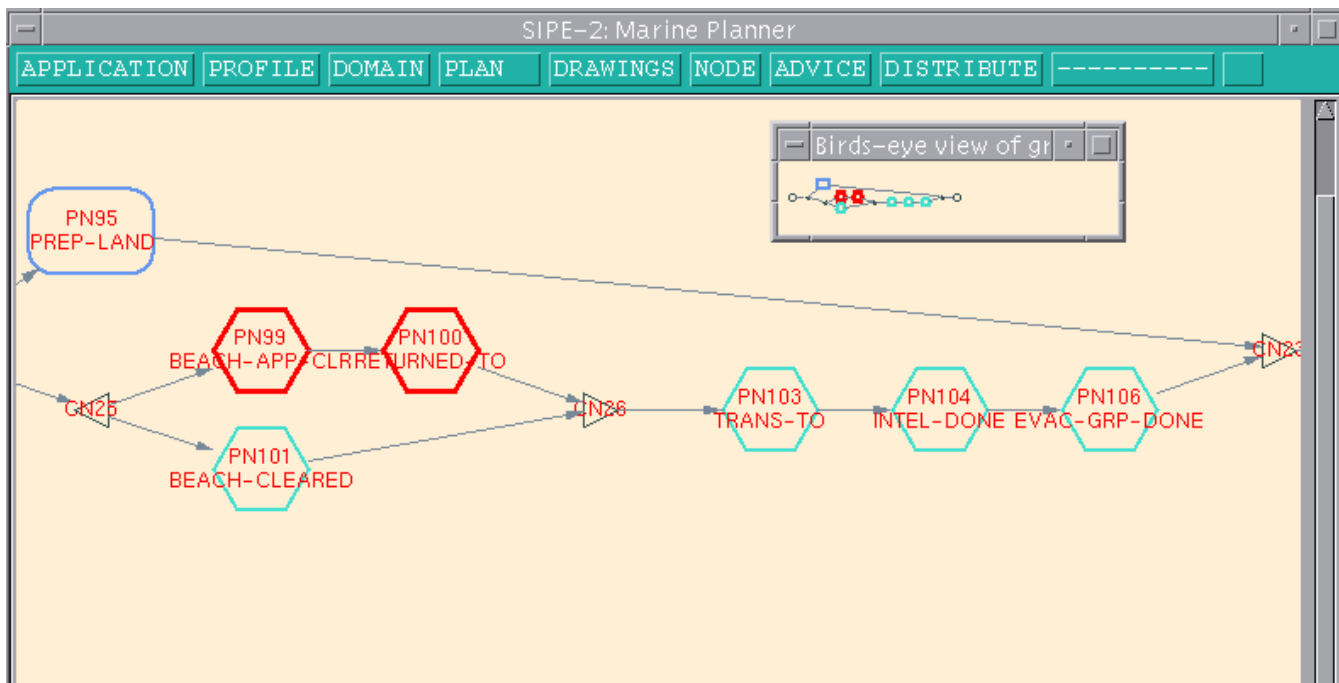


Figure 3: Marine Corps planning cell’s subplan before receiving Naval planning cell’s postconditions

identifies postconditions that may be relevant to other cells, as discussed in the next section. Suppose that cell A determines that a set of postconditions are relevant to remote cell B . Suppose further that these postconditions appear in the expansion of a high-level goal G in A ’s subplan. These postconditions are sent to cell B , which adds them to its subplan as additional nodes, appearing immediately before the placeholder node that corresponds to G . Along with the postconditions, ordering links and updated variable constraints are sent and attached to the added postcondition nodes.

The parts of a local subplan that a remote planning cell has been informed of make up a *skeletal subplan* that corresponds to the remote cell’s view into the local cell’s subplan. The local cell maintains a record of this skeletal subplan, i.e., it keeps track of the partial ordering associated with those postconditions that have been sent to the remote cell. The ordering links associated with the skeletal subplan are sent to the remote planning cell, so that it can construct the same partial ordering for those postconditions. When additional postconditions are sent to the remote cell, only the ordering links necessary to incrementally update the skeletal plan are sent along with the postconditions.

All of the constraints associated with variables in the postconditions are also passed to the remote agent. For example, if a postcondition is sent that states that unit X is located at point Y , the constraints on variables X and Y are also sent. These constraints can include class constraints (X is of type Force Reconnaissance Unit), properties (X has 40 troops), and instantiations (Y is bound to Pacifica-Beach-A).

The result is that each planning cell has a view of the overall plan that varies in detail: the view of the cell’s own subplan is complete, but the view of the parts of the plan being developed by other cells are more or less sketchy, depending on how much information is determined to be relevant.

Information Sharing

The goal of information sharing in DSIPE is to minimize message passing among planning cells (especially important in limited-bandwidth environments) while providing conflict detection and resolution early in the planning process. This is accomplished by sending constraints that arise during each cell’s planning process to other cells that may need to use those constraints, or that may generate subplans that conflict with them. Irrelevance reasoning is used to filter constraints that are guaranteed not to be of interest to other cells, thus reducing the communication requirements without sacrificing consistency.

Throughout the planning process, each planning system monitors the local cell’s planning activity for constraints and subgoals that may be relevant to other planning cells, and notifies those cells of this information. For example, the Naval planning cell might notify the Marine Corps planner that a particular landing area will be swept of mines by a specified time. Currently, the only constraints that are monitored in this way are the postconditions. In future work, we will develop methods to monitor temporal and other variables, preconditions, and resource constraints.

Relevant preconditions are those conditions that the local (sending) cell needs (or would like) to have the remote cell maintain or establish. For example, sending a precondition could tell the remote planning cell not to move a needed unit, or to perform a surveillance activity for an intelligence estimate.

We have implemented an algorithm for automatically filtering the postconditions that are sent from one planning cell to another by eliminating the postconditions that are provably irrelevant to the latter cell's decision-making process. The algorithm constructs a query tree (Levy & Sagiv 1993) for the objectives that have been assigned to a remote planning cell.³ This tree can be constructed by the local cell if planning operators are shared by the cells, or can be constructed remotely and sent to the local cell if each cell has its own set of planning operators.

The query tree identifies which predicates (i.e., which types of knowledge) are used in the preconditions of the operators that may be applied to expand those objectives. The postconditions within the generated subplan are matched to the predicates in this tree to determine whether the postcondition may be relevant to the remote planner (i.e., whether it potentially matches a precondition in the remote subplan).

Figure 4 shows a query tree for the BEACH-CLEARED goal in the maritime campaign planning example. The query tree is generated by identifying operators that could be used to achieve the BEACH-CLEARED goal, propagating any variable class constraints down the tree, and applying the process recursively to goals identified within the operators in the tree. The preconditions associated with the operators in the resulting tree, together with the class constraints, describe the classes of ground facts that are potentially relevant to achieving the top-level objective(s) in the query tree.

As reported in (Wolverton & desJardins 1998), irrelevance-based filtering can reduce message traffic substantially. For the two domains in which we ran experiments—a logistics domain and the maritime campaign planning domain presented here—using irrelevance-based filtering reduced the number of messages sent by 18% and 94%, respectively. Planning time was also reduced: since fewer remote postconditions are added to the local subplan, less computation needs to be done to compute the impact of these postconditions. Overall planning time was 35% lower in the logistics domain and 5% lower in the maritime planning domain.

During its planning process, the Naval planning cell (not shown) identifies some potentially relevant postconditions in its subplan. In particular, it finds that

³The query tree is similar to an *operator graph* (Smith & Peot 1993). Previous planning systems used operator graphs to identify threats that can be deferred (Smith & Peot 1993) or to further constrain the applicability of operators (Gerevini & Schubert 1996). In contrast, DSIPE uses the query tree to identify irrelevant classes of ground facts and to prevent them from being considered by the planner.

some of the goals and tasks that have been created during the expansion of its BEACH-APP-CLR subgoal may be relevant to the subgoals for which the Marine Corps planning cell is responsible. These postconditions correspond to the launching and movement of a small craft that the Navy will use to clear the approach to the beach. It sends these postconditions, along with ordering information that informs the Marine Corps planning cell of the order in which the postconditions should appear. Figure 5 shows the Marine Corps planner after these postconditions have been received and inserted into its subplan. The postconditions are encoded in the dark rounded rectangles that appear before the BEACH-APP-CLR goal.

After the Marine Corps cell has expanded its subplan further, Figure 6 shows that several dependencies, constraining the partial order of the subplan, have been introduced. These ordering constraints represent dependencies between the Naval planning cell's postconditions and tasks in the Marine Corps subplan. For example, the Marine Corps has determined that it can use the Navy's small craft to transport a group of Marines to the beach. The Marine Corps subgoal to launch a craft (CRAFT-TO, PM420) can be solved by using the Navy's craft; the Naval launch (CRAFT-TO, PM458) is then constrained to occur before the point in the plan where the Marine Corps requires the craft. The Marine Corps then introduces an additional RENDEZ action earlier in the plan to rendezvous with the Navy's craft for loading. With this dependency in place, it completes its planning, resulting in a subplan that is conditioned on the Navy's planned activities.

Plan Merging

After all of the subplans are complete, they are submitted to the coordinating planning cell and merged. The shared plan structure (corresponding to the skeletal plans in each cell) guides the merging process, resulting in a merged partially ordered network (see Figure 2). SIPE-2's planning critics are then run to identify conflicts that were not previously detected.

Because the relevant constraints have been shared during the planning process, the expectation is that few if any conflicts will appear during plan merging. However, because of the complexity of planning dependencies, conflicts can arise. For example, although the temporal constraints between subplans are shared, propagation of temporal constraints among subplans is performed only when postconditions are passed from one cell to another. The result is that cycles of temporal constraints that are distributed across subplans may not be detected until plan merging time.

If conflicts are detected during plan merging, in the current system, the coordinating cell resolves the conflict internally (e.g., by repairing the plan or by reallocating resources). In future work, we will incorporate methods for the coordinating cell to notify the subplanners to repair the conflict (see *Future Work—Conflict Resolution*).

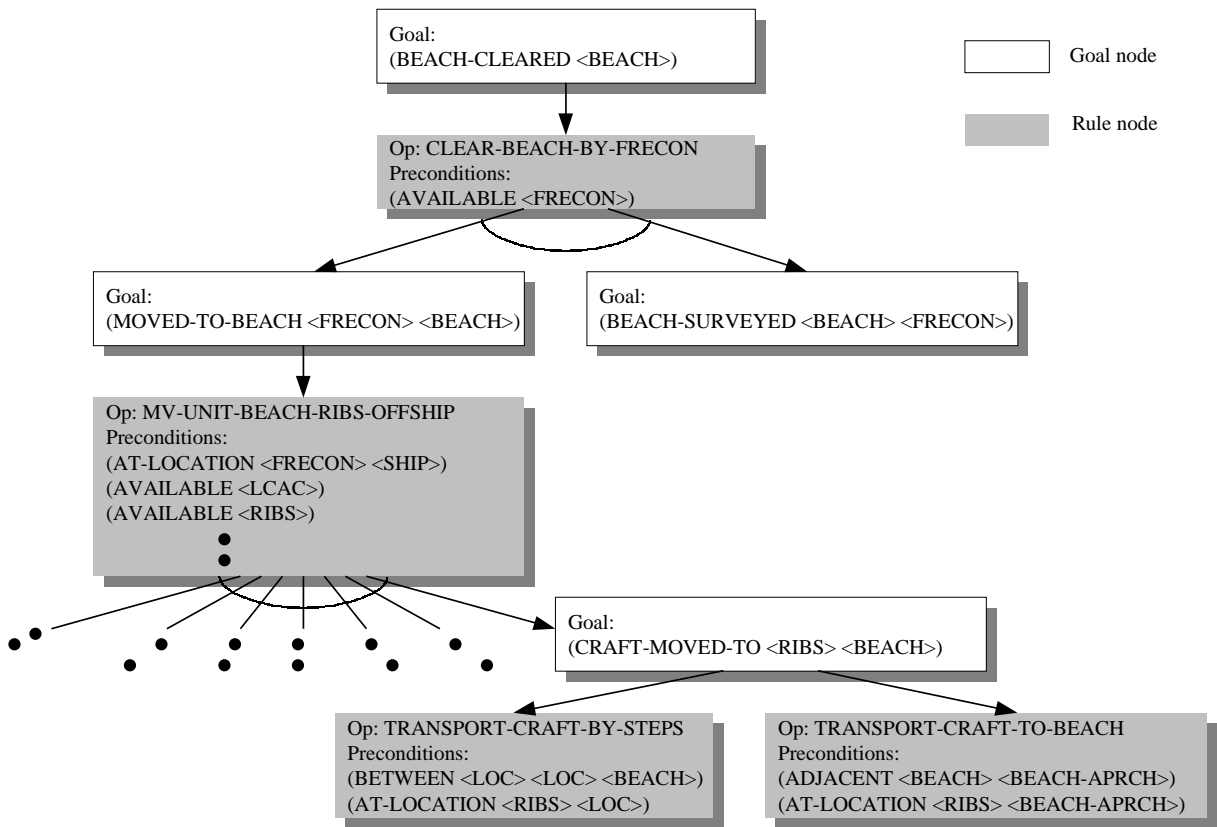


Figure 4: Query tree for the BEACH-CLEARED goal

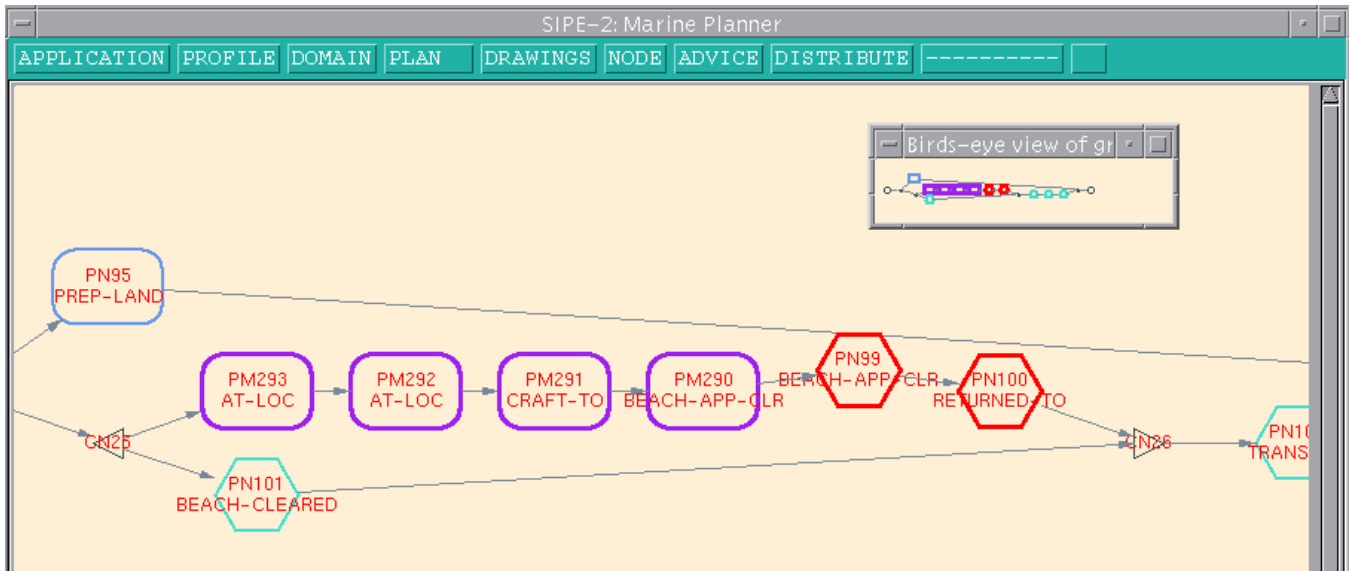


Figure 5: Marine Corps planning cell's subplan after receiving relevant postconditions (dark rounded rectangles) from the Naval planning cell

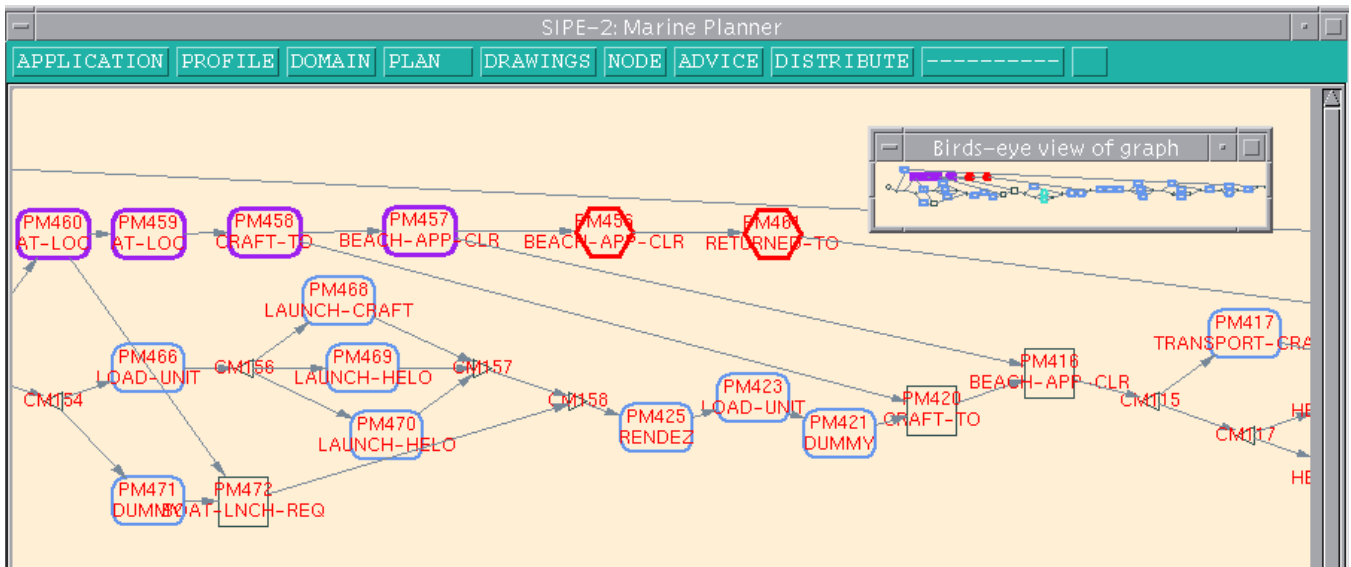


Figure 6: Interdependencies between the Marine Corps and Naval subplans, from the viewpoint of the Marine Corps planning cell

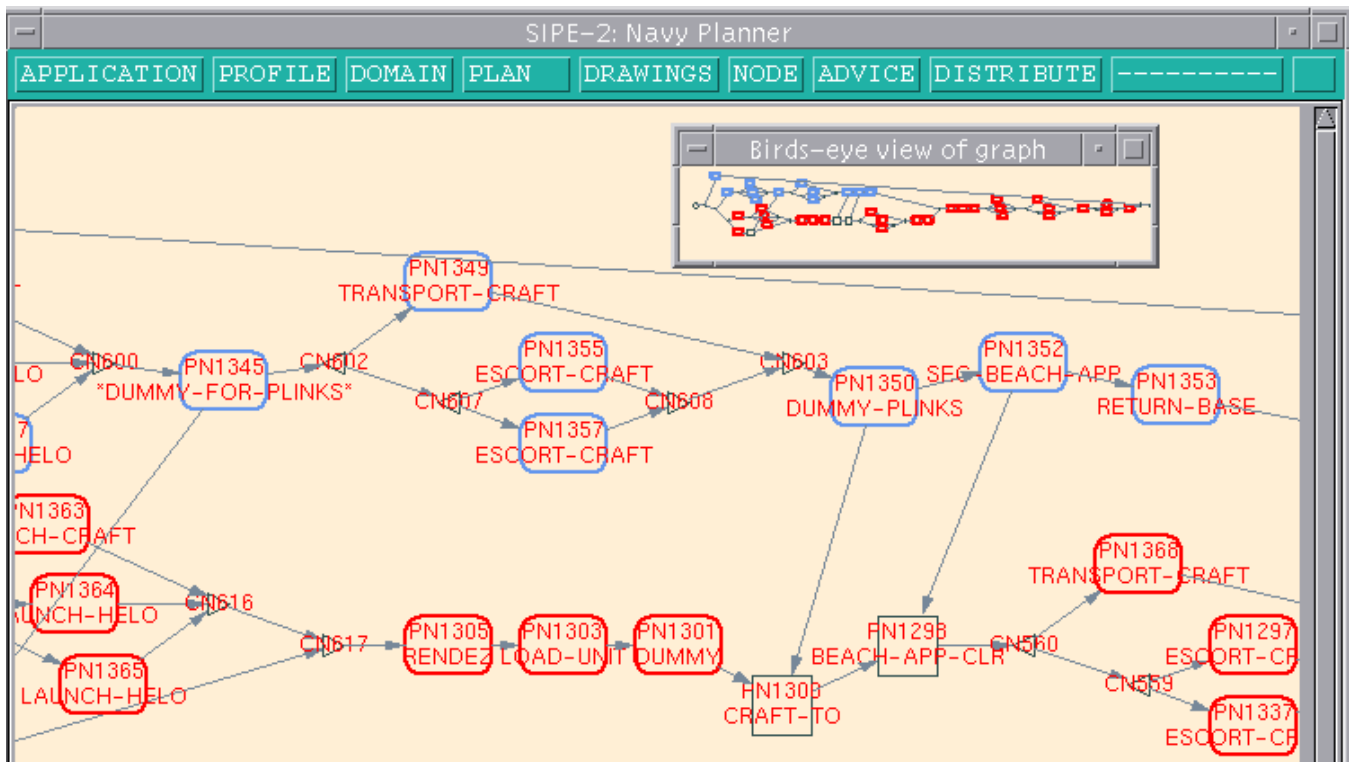


Figure 7: Final merged plan as viewed from the coordinating planning cell

In the maritime planning example, the subplans are now complete, and are submitted to the coordinating planning cell, which merges them. (In this case, the Naval planning cell is serving as the coordinating planning cell as well, but in the general case, the coordinating planning cell could be distinct from the other planning cells.) Figure 7 shows the merged plan. Although this display does not show the complete detailed plan, several parallel interactions can be seen, indicated by the three links from the upper part of the plan to the lower parallel branch. These interactions occur between the Marine Corps and Naval subplans, and are all identified during the distributed planning process, then maintained when the plans are merged.

Future Work

The primary research directions that we are exploring are synchronization of planning cells, conflict resolution and negotiation, and improved information management. These are discussed in the following subsections.

Synchronization of Distributed Planners

One of the central research issues in DSIPE is how to communicate effectively about the planning process so that distributed planning cells can generate subplans on independent timelines, detect and resolve potential conflicts early, avoid duplicating effort, and create a high-quality merged plan at the end of the process.

The current distributed planning process is somewhat inflexible, in that the primary way that planning cells can share knowledge is via communication of postconditions that are created in one of the subplans. Also, variable constraints are only posted when potentially relevant postconditions are generated. Therefore, if one planning cell's preconditions depend on the postconditions of another cell, the part of the latter's subplan that produces the necessary postconditions must be completed first, so that the former cell can use the postconditions in constructing its subplan.

We are currently developing a more flexible process that will update information about variable constraints and instantiations dynamically as they change during planning. It will also allow other types of distributed constraints to be posted—for example, one cell could inform another cell of preconditions that it expected to be maintained or created in the world, or could explicitly ask for assistance in achieving an objective.

We are also extending DSIPE's search control methods so that subplans can be generated in different ways, depending on the situation. For example, a planning cell might decide to delay achievement of a particular objective (while continuing to plan out other objectives), pending more information about other cells' subplans. In order to support this flexible search control, we will explore ways in which a planning cell can model remote subplans and planning processes in more detail.

Conflict Resolution

To date, we have focused on methods for representing and managing the distributed planning process, and for identifying constraint violations. However, making this technology robust and applicable to a broad range of distributed planning problems will also require developing techniques for resolving these conflicts.

The approach we are taking is to involve the user in the decision-making process, while automating as much as possible those details of conflict detection and resolution that do not require human intervention. When a conflict is detected (or a potential conflict is predicted), either during distributed planning or at plan merging time, the system will automatically analyze the conflict and determine a set of appropriate conflict resolution strategies. These strategies may include replanning one or both subplans, agreeing to violate a precondition, requesting additional resources from the coordinating planning cell, or relaxing soft constraints such as arrival time of a unit. These will be ranked and presented to the user, with information about the tradeoffs of the various strategies, in order to facilitate their decision-making. If the conflict occurs during plan merging, additional heuristics will be invoked to decide whether the coordinating planning cell should repair the conflict on its own, or whether the planning cells that generated the subplans involved in the conflict should collaboratively repair the conflict.

We are also exploring methods of opportunistic plan improvement where no conflict exists—for example, identifying and removing redundant parts of the plan, or replanning to increase cooperation among the planning cells.

Information Management

The irrelevance-based filtering technique has been shown to reduce message traffic in DSIPE. The current implementation considers only the postconditions that are generated within each agent's subplan. The technique is general, however, and could be applied to other types of message traffic. For example, irrelevance reasoning could be used to determine whether to send requests for other agents to solve subgoals, or requests to achieve or maintain preconditions. Another application of irrelevance reasoning would be to use the query trees to allocate subgoals among agents, based on some notion of locality, in order to minimize interactions among subplans. This approach is similar to how COLLAGE (Lansky 1994; Lansky & Getoor 1995) uses *localization* to decompose a planning problem.

Query trees could be dynamically pruned as subplans are generated (e.g., when an alternative planning strategy has been ruled out, the associated preconditions would no longer be relevant). This would entail notifying remote cells of updates to the tree, so there would be a tradeoff between the bandwidth required for the updates and the reduction in message traffic that would result from using the updated trees.

The irrelevance reasoning approach could also be extended by incorporating notions of probability and utility: how *likely* is it that an agent will in fact need to know a piece of information, and how *useful* will it be to the agent to know it? STEAM (Tambe 1997), a distributed agent architecture that extends joint intentions theory (Levesque, Cohen, & Nunes 1990) to maintain a coherent view of the team's goals and plans, uses a decision-theoretic framework that incorporates the costs and benefits of communication, as well as the probability that other agents already have the information in question. This approach, which is also used in (Gmytrasiewicz, Durfee, & Wehe 1991), could be layered on top of DSIPE's irrelevance-based filtering to incorporate additional decision-making factors in deciding whether and when to communicate information between planning agents.

Conclusions

We have described DSIPE, an implemented system that extends the state of the art of distributed HTN planning. DSIPE provides decision support to human planners in a collaborative planning environment. The key contributions of this research are a shared plan representation, methods for automatically identifying and sharing relevant information among planners, and plan merging techniques. We have also discussed important future research directions, including flexible synchronization of distributed planners, conflict negotiation and resolution methods, and system evaluation.

Acknowledgments

This research was funded by the Office of Naval Research (ONR) and the Space and Naval Warfare Systems Command (SPAWAR) under contract N66001-97-C-8515. The authors would like to thank Mr. Mike Gherrity and Mr. Dave Swanson of SPAWAR for their guidance. Thanks also to David Wilkins for consulting and implementation assistance on SIPE-2 issues.

References

- Corkill, D. D. 1979. Hierarchical planning in a distributed environment. In *Proceedings of the 6th International Joint Conference on Artificial Intelligence*. San Francisco: Morgan Kaufmann.
- Finin, T.; Labrou, Y.; and Mayfield, J. 1997. KQML as an agent communication language. In Bradshaw, J., ed., *Software Agents*. Cambridge, MA: MIT Press.
- Gerevini, A., and Schubert, L. 1996. Accelerating partial-order planners: Some techniques for effective search control and pruning. *Journal of Artificial Intelligence Research* 5:95–137.
- Gmytrasiewicz, P. J.; Durfee, E. H.; and Wehe, D. K. 1991. The utility of communication in coordinating intelligent agents. In *Proceedings of the 9th National Conference on Artificial Intelligence*, 166–172. Menlo Park, CA: AAAI Press.

- Janssen, B.; Spreitzer, M.; Larner, D.; and Jacobi, C. 1998. ILU 2.0alpha12 reference manual. Technical report, Xerox PARC, Palo Alto, CA.
- Lansky, A. L., and Getoor, L. C. 1995. Scope and abstraction: Two criteria for localized planning. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, 1612–1619. San Francisco: Morgan Kaufmann.
- Lansky, A. L. 1994. Located planning with diverse plan construction methods. Technical Report FIA-TR-9405, NASA Ames Research Center, Mountain View, CA.
- Levesque, H. J.; Cohen, P. R.; and Nunes, J. H. T. 1990. On acting together. In *Proceedings of the 8th National Conference on Artificial Intelligence*, 94–99. Menlo Park, CA: AAAI Press.
- Levy, A. Y., and Sagiv, Y. 1993. Exploiting irrelevance reasoning to guide problem solving. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence*. San Francisco: Morgan Kaufmann.
- Smith, D. E., and Peot, M. A. 1993. Postponing threats in partial-order planning. In *Proceedings of the 11th National Conference on Artificial Intelligence*, 500–506. Menlo Park, CA: AAAI Press.
- Tambe, M. 1997. Agent architectures for flexible, practical teamwork. In *Proceedings of the 14th National Conference on Artificial Intelligence*. Menlo Park, CA: AAAI Press.
- Wilkins, D. E. 1988. *Practical Planning: Extending the Classical AI Planning Paradigm*. San Francisco: Morgan Kaufmann.
- Wolverton, M. J., and desJardins, M. 1998. Controlling communication in distributed planning using irrelevance reasoning. In *Proceedings of the 15th National Conference on Artificial Intelligence (AAAI-98)*, 868–874. Menlo Park, CA: AAAI Press.