Performance of Coordinating Concurrent Hierarchical Planning Agents Using Summary Information

Bradley J. Clement and Edmund H. Durfee
University of Michigan
Artificial Intelligence Laboratory
1101 Beal Avenue, Ann Arbor, MI 48109-2110, USA
{bradc, durfee}@umich.edu

In a shared environment with limited resources, an agent may have enough information about the environment to individually plan its course of action but may not be able to anticipate how the actions of others will interfere with accomplishing its goals. Prior techniques have enabled such agents to cooperatively seek merges of individual plans that will accomplish all of their goals if possible [4]. In addition, recent research has provided these agents with tools to coordinate their *concurrent hierarchical plans* (CHiPs) resulting in more flexible abstract solutions that allow the agents to choose refinements of their actions during execution that can withstand some amount of failure and uncertainty [2]. In order to reason about abstract plans to identify and resolve conflicts, information about how an abstract plan must or may be refined into lower level actions must be available. This information can be summarized from the conditions of the abstract plan' subplans in its potential refinements. Based on this summary information, methods have been developed to soundly and completely test whether particular ordering constraints will or might resolve all of the conflicts among the plan hierarchies [1].

We can integrate these tests into an algorithm that a group of agents can use to search for a consistent coordinated plan. This algorithm is sound and complete, returns the optimal global plan if it exists, and uses heuristics that take advantage of summary information. A coordinating agent first collects the summary information of the top-level plans of each agent whose plans may conflict. The coordinator tries to find an ordering of the plans that resolves all conflicts, and the agents or the coordinator then chooses subplans of the global plan to expand deeper and deeper (by transferring summary information of the more detailed agents' subplans to the coordinator) until the optimal solution is found or the search space has been exhausted. A state of the search is a partially elaborated global plan that we represent as a conjunction of frontier plan steps for all of the agents, a set of temporal constraints, and a set of blocked plans. The set of temporal constraints includes synchronization constraints added during the search in addition to those dictated by the agents' individual hierarchical plans. The blocked subplans are the pruned decomposition choices.

This work was supported in part by NSF (IRI-9158473) and DARPA (F30602-98-2-0142).

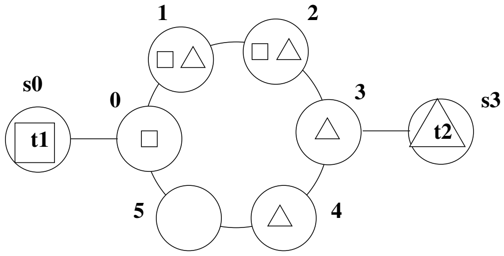

Figure 1. Transports $t1$ and $t2$ must pick up square and triangle evacuees respectively and return them to their original safe locations without colliding on single lane routes.

The operators of the search are expanding non-primitive plans, selecting subplan choices, and adding temporal constraints on pairs of plan steps. When a plan is expanded, it is replaced by its subplans, and the ordering information is updated in the global plan. In our experiments, we separate the search for synchronizations from the expansion and selection of subplans. An outer search explores the space of plans at different levels of abstraction. For each state in the outer search, an inner search explores the space of synchronizations by resolving threats with ordering constraints. For example, in the simple evacuation domain shown in Figure 1, agents $t1$ and $t2$ can use their summary information of just their highest level plans (whose abstract goals are to perform the evacuation) to determine that if $t1$ executes its plan completely before $t2$ starts, then they can avoid a collision and accomplish their goals no matter how they traverse the ring. However, if they wish to act concurrently, then they must expand their hierarchies deeper, block subplans that result in unavoidable collisions, and add ordering constraints to lower level plan steps in the hierarchy such that their movements are synchronized to avoid collisions.

Summary information can also be valuable in directing the search to avoid branches in the search space that lead to inconsistent or suboptimal global plans. Inconsistent global plans can be pruned away at the abstract level by doing a quick check to see whether there are irresolvable conflicts. In terms of the number of states expanded during the search, employing this technique will always do at least as well as not using it. Another strategy that is employed is to first expand plans involved in the most threats. This "expand on most threats first" (EMTF) heuristic aims at driving the search down through the hierarchy to find the subplan(s) causing conflicts with others so that they can be resolved more quickly. Another heuristic used in parallel in our experiments is "fewest threats first" (FTF). Here the search orders nodes in the outer search queue by ascending numbers of threats to resolve. By trying to resolve the threats of global plans with fewer conflicts, it is hoped that solutions can be found more quickly. So, EMTF is a heuristic ordering plans to expand, and FTF orders subplan choices and, thus, search states to investigate. In addition, in trying to find optimal solutions in the style of a branch-and-bound search, we use the cost of abstract solutions to prune away branches of the search space whose minimum cost is greater than the maximum cost of the current best solution. Again, pruning abstract plans can only help improve the search.

Using summary information to find abstract solutions to the coordination problem can improve the overall performance of coordinating and executing plans [2]. Although there is anecdotal evidence that coordination is cheaper at higher levels in the hierarchy, even in the worst case where summary conditions increase exponentially up the hierarchy, finding solutions at abstract levels is expected to be exponentially cheaper than at lower levels.

At the primitive level d of a global plan hierarchy of depth d, each plan has only $3c' = O(c')$ conditions, and there are $c' + bc' = O(bc')$ summary conditions for each plan at level $d-1$ and $O(b^2c')$. There are at most $3nc'$, or $O(b^dc')$, summary conditions at the root of the hierarchy—this is the total number of pre-, in-, and postconditions in the hierarchy. At the ith depth level in the hierarchy, each of the $O(b^i)$ plans has $O(b^{d-i}c'^2)$ summary conditions in the worst case. As described in [1], the algorithm to check whether a particular ordering of n plan steps (each with c summary conditions) results in all plans executing successfully is similar to deriving their collective summary information and has a complexity of $O(n^2c^2)$. Checking such a synchronization for the plans at any level i in a plan hierarchy is, thus, $O(b^{2i}b^{2d-2i}c'^2) = O(b^{2d}c'^2)$. So, since i drops out, the complexity of doing this check is *independent of the depth level*. But, there is a huge space of $n! = O((b^i)!)$ sequential orderings of the n plans at level i to potentially check to find a valid synchronization. Thus, the search space grows doubly exponentially down the hierarchy despite the worst case when the number of conditions grows exponentially up the hierarchy. We also prove that the problem of finding a valid synchronization in this space is intractable (NP-complete) by reducing HAMILTONIAN PATH to the THREAT RESOLUTION problem for STRIPS planning [3]. Thus, the improvements made using summary information can yield exponential savings while only incurring a small polynomial overhead in deriving and using summary information.

If the goal is to find the optimal solution, a reasonable criticism might be that using summary information to reason at abstract levels will be more costly than just coordinating at the lowest level of primitive actions because of the overhead of deriving and using summary information. Our experimental results contradict this criticism and show how the algorithm and heurisitics described above can better focus the search to much more quickly find optimal solutions at lower levels than a FAF ("fewest alternatives first") heuristic [5] that chooses to expand subplan choices that have the fewest number of subplans.

Plans were generated for an evacuation domain similar to the one in Figure 1. Agents had plans to either visit their specified locations by traveling in one direction only or switching directions at some location. These choices expand into choices to begin traveling clockwise or counterclockwise. For the branch where the agent switches directions, it can choose to change directions at any location it is specified to visit. Optimal solutions are coordinated plans with the shortest total completion time where each move has a uniform time cost. We chose problems with four, six, and eight locations; with two and three agents; and with no, some, and complete overlap in the locations the agents visited in order to vary the amount of plan interactions.

For problems with only four locations and two agents, both algorithms found optimal solutions quickly. For more complex problems, the heuristics using summary information appear to make great improvements over FAF, which could only solve six of the 21 problems within memory constraints. For the problems where it did find solutions, FAF required an average of 2.5 orders of magnitude more CPU time. These results are by no means conclusive, but they do show promise for search based on summary information. For most of these problems, coordinating at the primitive level was intractable, but the algorithm using summary information was able to find an abstract solution quickly.

We have shown that summary information can find solutions at higher levels exponentially more quickly than at lower levels; and we have identified heuristics and search techniques that can take advantage of summary information in finding coordinated plans. In addition, we have characterized a coordination algorithm that takes advantage of these search techniques and experimentally shown how it can make large improvements over an FAF heuristic in finding optimal coordinated plans. More work is needed to show that these results translate to different domains, and future considerations include comparing this approach to other planning heuristics that capitalize on domain knowledge in order to better understand the relationship between plan structure and search performance. We expect the benefits of using summary information to also apply to hierarchical planning and wish to compare these techniques with current heuristics for concurrent hierarchical planning.

References

[1] B. Clement and E. Durfee. Theory for coordinating concurrent hierarchical planning agents. In *Proc. AAAI*, 1999.

[2] B. Clement and E. Durfee. Top-down search for coordinating the hierarchical plans of multiple agents. In *Proc. Intl. Conf. Autonomous Agents*, 1999.

[3] B. Clement and E. Durfee. Theory for coordinating concurrent hierarchical planning agents. http://www.eecs.umich.edu/~bradc/papers/extrep.ps, 2000.

[4] M. P. Georgeff. Communication and interaction in multiagent planning. In *Proc. AAAI*, pages 125–129, 1983.

[5] R. Tsuneto, J. Hendler, and D. Nau. Space-size minimizationin refinement planning. In *Proc. Fourth European Conference on Planning*, 1997.