Task Fusion Heuristics for Coalition Formation and Planning

Robotics Track

Gilberto Marcon dos Santos
Oregon State University
marcondg@oregonstate.edu

Julie A. Adams
Oregon State University
julie.a.adams@oregonstate.edu

ABSTRACT

Automating planning for large teams of heterogeneous robots is a growing challenge. The planning literature incorporates expressive features, but examples that scale to multiple robots in complex domains are limited and fail to generate feasible plans. The Coalition Formation then Planning framework accelerates planning by decomposing the robots into coalitions, allocating tasks to each coalition, and planning each task separately. However, the task decomposition limits cooperation between coalitions and results in many nonexecutable plans. The presented Task Fusion heuristics fuse coalition-task pairs, resulting in higher success rates by leveraging relaxed plans to estimate couplings between tasks and determine the coalition-task pairs to be fused. The heuristics are compared to baseline methods across randomly generated problems that incorporate temporal and continuous constraints.

KEYWORDS

coalition formation; temporal planning; continuous planning

1 INTRODUCTION AND BACKGROUND

Robots are rapidly moving into the commercial, medical, and military domains. Robots have proven potential to assist in first response to major disasters, such as search and rescue, bomb defusal, and firefighting, but currently require micromanagement [1]. Decision making is limited to low level actions. Exploiting the full potential of autonomous robots requires scalable automated planning capable of accurately modeling the dynamic and uncertain real-world problems that incorporate a diverse set of robots [1].

First response requires rapid evaluation and deployment of available personnel and equipment to mitigate the situation. As technological capabilities improve, the complexity of the deployment allocation and assignment problem increases. Existing planning methods (e.g., [2–4, 7, 10–12, 14–16]) fail to account for all of the domain's complexities, such as requiring continuous fluents, concurrent actions, and real-time results that are robust to dynamic and uncertain situations. Existing planning methods partially meet the requirements, but cannot scale to dozens of robots [8].

Dukeman [5] developed the Coalition Formation then Planning (CFP) framework for merging automated planning and coalition formation to improve scalability when developing plans for complex missions composed of dozens of robots for domains that have continuous and temporal requirements. CFP assigns robots to coalitions according to their capabilities, and allocates tasks to each coalition. Planning for each task separately dramatically accelerates planning, but limits cooperation between coalitions. Task Fusion (TF) extends CFP by merging coalitions based on evaluating the utility of fusing each coalition pair [5]. TF heuristics estimate utility and fuse the highest scoring pairs; thus, allowing explicit cooperation between robots and improving plan success rate. However, the TF effectiveness was inconclusive [5]. The existing heuristics [5] estimate the utility of fusing coalition-task pairs, p_i and p_j.

The *Coalition Similarity* (CS) heuristic [5] operates on coalition-task pairs that share common agents. Coalition-task pairs with no common agents score 0, while those that share all agents score 1. The *Coalition Assistance* (CA) heuristic [5] estimates the ratio of coalition capabilities over task requirement capabilities after fusion, prioritizing pairs that share the same task requirement capabilities [6]. These heuristics do not consider planning-related information, such as robots handling the same logical objects or sharing the same physical room, which limits the heuristic's accuracy and leads to poor success rates. Task coupling can be an accurate estimate of Task Fusion (TF) utility. Detecting task couplings allows fusing coupled tasks and planning tasks together, which improves cooperation and produces higher plan success rates.

2 TASK FUSION HEURISTICS

A relaxed plan can be used as a TF heuristic to leverage the information revealed during planning [6]. Relaxed plans provide an estimate of the full plan, including sequences of actions to accomplish the plan, logical objects involved in each action, and the time at which each action occurs. The simplest approach leverages relaxed plans' contents for estimating TF utility to analyze the overlap of actions and logical objects between two relaxed plans. This overlap indicates how similar the plans are and how coupled each plans' tasks are, which can reduce the length of the fused plan. Higher overlap of actions and logical objects for coupled tasks indicates the robots interact with common objects and navigate through common locations, which are also represented as logical objects. The planner will consider each tasks' goals individually, and potentially produce redundant sequences of actions, when two coupled tasks are planned separately Planning uncoupled tasks together does not improve plan success rates, and often increases planning time and memory usage. Uncoupled tasks can be planned and executed separately; thus, the utility of fusing uncoupled tasks is lower than that of fusing coupled tasks. The proposed TF heuristics use the overlap of actions and logical object occurrences to estimate task coupling and the utility of fusing coalition-task pairs.

A relaxed plan, π^r, consists of a list of actions, where each action entry contains a start time t, robots $A = \{A_1, A_2, \ldots\}$, and planning-model first-order logic objects $o = \{o_1, o_2, \ldots\}$. Relaxed

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10–15, 2018, Stockholm, Sweden

plan heuristics compile a list of logical object and action occurrences, extracted from each relaxed plan action entry. Each action-object occurrence, tagged with the associated action start time, τ, populates the action-object list, $L = \{\langle o_1, \tau_1\rangle, \langle o_2, \tau_2\rangle, \ldots\}$. The similarities between plans π_i^r and π_j^r render an estimate for the utility of fusing coalition-task pairs p_i and p_j.

Let π_i^r and π_j^r represent the relaxed plans for coalition-task pairs p_i and p_j, respectively. Relaxed plan heuristics are a function $H(\pi_i^r, \pi_j^r) : \pi_i^r \times \pi_j^r \to \mathbb{R}^+$, that maps to a utility value and require synthesizing relaxed plans π^r for all m coalition-task pairs p, but leverages plans' details that are otherwise unavailable.

The *Object* (O), *Action* (A), and *Action-Object* (AO) heuristics represent the level of overlap between the object and action occurrences in relaxed plans π_i^r and π_j^r for coalition-task pairs p_i and p_j, respectively: $\frac{1}{|L_i| \cdot |L_j|} \cdot \sum_{l_i \in L_i} \sum_{l_j \in L_j} (l_i = l_j)$, where $|L_i|$ and $|L_j|$ are list sizes for action-object lists L_i and L_j, respectively. All pairs of entries from both action-object lists are compared. Each heuristic variant populates the relaxed plan lists, L_i and L_j, with object occurrences, action occurrences, or both. The *Object heuristic* populates lists with object occurrences; the *Action heuristic* populates lists with action occurrences; and the *Action-Object heuristic* populates lists with both action and object occurrences. The normalizing fraction ensures that the heuristic values are between $[0, 1]$, where 0 represents no task coupling and 1 indicates maximal task coupling.

The *Object-Temporal* (OT), *Action-Temporal* (AT), and *Action-Object-Temporal* (AOT) heuristics integrate temporal dependencies in order to account for action and object interactions at different times throughout the plan. The temporal variants populate the relaxed plan lists with object and action occurrences and weight each matching list entry with a decaying exponential weighting factor. The weighting ranks pairs that interact with the same objects at similar times higher than pairs that interact with the same objects at different times. The weighting factor is a function of the time difference between each matching list entry: $\frac{1}{|L_i| \cdot |L_j|} \cdot \sum_{l_i \in L_i} \sum_{l_j \in L_j} (l_i = l_j) \cdot e^{-|\tau_i - \tau_j|}$, where τ_i and τ_j are temporal timestamps for list entries l_i and l_j, respectively. The weighting factor is 1 if $\Delta\tau = |\tau_i - \tau_j| = 0$, (i.e., the object matching occurs at the same time), and 0 if $\Delta\tau \to \infty$, (i.e., the object matching occurs at different times).

3 EXPERIMENTS AND RESULTS

The planning outcomes were: Success, a valid plan; Nonexecutable, no plan can be derived given the coalition's composition and allocated tasks; Time Fail, the time limit was exceeded; and Memory Fail, the memory limit was exceeded. Ten robot coalitions and ten missions were randomly generated to form 100 problems per domain. The Blocks World Domain [9] incorporated temporal constraints, continuous fluents and modeled a variety of end-effectors, block sizes, and multiple robot arms. The First Response Domain [6] models disaster response problems that require coordinating heterogeneous human-robot teams and incorporates time-varying fluents, so that the robot batteries drain as a function of robot activity over time. Human-robot teams cooperate to rescue victims, collect hazardous objects, clear gas leaks, and clear blocked roads after a natural disaster.

RACHNA [17], a market-based coalition formation algorithm, was used for the First Response Domain and a dynamic programming coalition formation algorithm [13] was used for the Blocks World Domain. Both the regular and the relaxed plans were synthesized using the COLIN package [4]. Each trial was time capped at one hour and memory usage was limited to 120 GB.

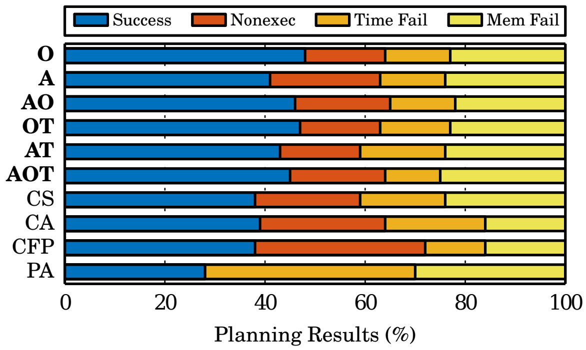

Figure 1: Blocks World Domain planning results.

Blocks World Domain: The Object (48%), Object-Temporal (47%), and the Action-Object (46%) heuristics produced the best planning success rates, as presented in Figure 1. Planning Alone had zero nonexecutable coalitions, but had the worst planning success (28%), time failure (42%), and memory failure (30%) rates. CFP produced the most nonexecutable coalitions (34%). The relaxed plan heuristics provide higher plan success rates; thus, larger temporal continuous planning problems involving multiple robots can be solved.

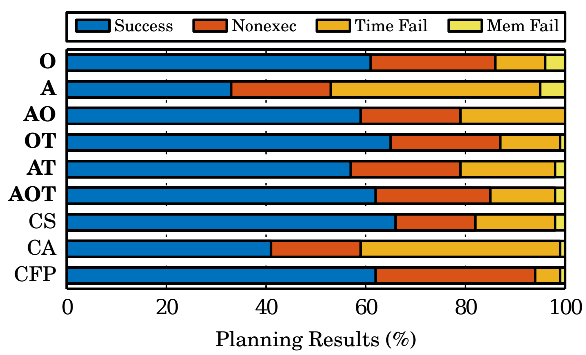

Figure 2: First Response Domain planning results.

First Response Domain: The Coalition Similarity heuristic produced the best planning success rate (66%), as presented in Figure 2, while the Object-Temporal heuristic was the second best (65%). Planning Alone exceeded the 1-hour processing time limit for all problems. CFP produced the highest rate of nonexecutable coalitions (32%), while the Action heuristic produced the highest failure rates, due to exceeding both the 1-hour time limit (42%) and the 120 GB memory limit (5%).

ACKNOWLEDGMENTS

This work was partially supported by NSF grant #1723924.

REFERENCES

[1] Ron Alterovitz, Sven Koenig, and Maxim Likhachev. 2016. Robot planning in the real world: Research challenges and opportunities. *AI Magazine* 2 (2016), 76–84.

[2] Daniel Bryce, Sicun Gao, David J. Musliner, and Robert P. Goldman. 2015. SMT-based nonlinear PDDL+ planning. In *Proceedings of the AAAI Conference on Artificial Intelligence*. 3247–3253.

[3] Amanda J. Coles, Andrew I. Coles, Maria Fox, and Derek Long. 2011. POPF2: A forward-chaining partial order planner. In *The International Planning Competition*. 65–70.

[4] Amanda J. Coles, Andrew I. Coles, Maria Fox, and Derek Long. 2012. COLIN: Planning with continuous linear numeric change. *Journal of Artificial Intelligence Research* (2012), 1–96.

[5] Anton Dukeman. 2017. *Hybrid mission planning with coalition formation*. Ph.D. Dissertation. Vanderbilt University, Nashville, Tenessee.

[6] Anton Dukeman and Julie A. Adams. 2017. Hybrid mission planning with coalition formation. *Journal of Autonomous Agents and Multi-Agent Systems* 31, 6 (Nov. 2017), 1424–1466.

[7] Patrick Eyerich, Robert Mattmüller, and Gabriele Röger. 2012. Using the context-enhanced additive heuristic for temporal and numeric planning. In *Towards Service Robots for Everyday Environments*. Springer, 49–64.

[8] Malik Ghallab, Dana Nau, and Paolo Traverso. 2016. *Automated planning and acting*. Cambridge University Press.

[9] Naresh Gupta and Dana S. Nau. 1992. On the complexity of blocks-world planning. *Artificial Intelligence* 56, 2-3 (1992), 223–254.

[10] Thomas Keller and Patrick Eyerich. 2012. PROST: Probabilistic planning based on UCT. In *Proceedings of the International Conference on Automated Planning and Scheduling*. 119–127.

[11] Hanna Kurniawati, David Hsu, and Wee S. Lee. 2008. SARSOP: Efficient point-based POMDP planning by approximating optimally reachable belief spaces. In *Proceedings of Robotics: Science and Systems*. 65–72.

[12] Masood F. Rankooh and Gholamreza Ghassem-Sani. 2015. ITSAT: An efficient SAT-based temporal planner. *Journal of Artificial Intelligence Research* (2015), 541–632.

[13] Travis C. Service and Julie A. Adams. 2011. Coalition formation for task allocation: Theory and algorithms. *Autonomous Agents and Multi-Agent Systems* 2 (2011), 225–248.

[14] Adhiraj Somani, Nan Ye, David Hsu, and Wee S. Lee. 2013. DESPOT: Online POMDP planning with regularization. *Advances in Neural Information Processing Systems* (2013), 1–9.

[15] Alvaro Torralba, Vidal Alcazar, Daniel Borrajo, Peter Kissmann, and Stefan Edelkamp. 2014. SymBA*: A symbolic bidirectional A* planner. In *The International Planning Competition*. 105–108.

[16] Vincent Vidal. 2004. The YAHSP planning system: Forward heuristic search with lookahead plans analysis. In *The International Planning Competition*. 56–58.

[17] Lovekesh Vig and Julie A. Adams. 2006. Market-based multi-robot coalition formation. In *Distributed Autonomous Robotic Systems 7*, Maria Gini and Richard Voyles (Eds.). Springer Japan, Tokyo, 227–236.