Effectively Incorporating Weighted Cost-to-go Heuristic in Suboptimal CBS (Extended Abstract)

Rishi Veerapaneni, Tushar Kusnur, Maxim Likhachev

Robotics Institute, Carnegie Mellon University
{rveerapa, tkusnur, mlikhach}@andrew.cmu.edu

Abstract

Conflict-Based Search (CBS) is a popular multi-agent path finding (MAPF) solver that employs a low-level single agent planner and a high-level constraint tree to resolve conflicts. The majority of modern MAPF solvers focus on improving CBS by reducing the size of this tree through various strategies with few methods modifying the low level planner. All low level planners in existing CBS methods use an unweighted cost-to-go heuristic, with suboptimal CBS methods also using a conflict heuristic to help the high level search. Contrary to prevailing beliefs, we show that the cost-to-go heuristic can be used significantly more effectively by weighting it in a specific manner alongside the conflict heuristic. We introduce two variants of doing so and demonstrate that this change can lead to 2-100x speedups in certain scenarios. Additionally, we show the first theoretical relation of prioritized planning and bounded suboptimal CBS and demonstrate that our methods are their natural generalization.

Introduction

Multi-Agent Path Finding (MAPF) is the problem of computing collision-free paths for a team of agents in a known environment while minimizing a measure of their travel times. This is required for several real-world tasks such as the smooth operation of automated warehouses (Li et al. 2020b), robot soccer (Biswas et al. 2014), coverage (Kusnur et al. 2021), and others.

Prioritized Planning (PP) (Erdmann and Lozano-Perez 1987) is a fast multi-agent planning approach that sequentially plans agents avoiding earlier "more important" agents, and has been applied to several domains (Čáp et al. 2015; Velagapudi, Sycara, and Scerri 2010). However PP provides no guarantees on completeness or suboptimality.

Conflict-Based Search (CBS) is a popular complete and optimal MAPF solver that employs a low-level single agent planner and a high-level constraint tree (CT) to resolve conflicts. Several methods speed up CBS by reducing the CT size by explicitly pruning branches, selectively expanding branches, adding sets of constraints, detecting symmetries, and improving high-level heuristics (Boyarski et al. 2015, 2021; Li et al. 2019, 2020a, 2021).

Enhanced CBS (ECBS) (Barer et al. 2014) introduced the first bounded-suboptimal version of CBS, utilizing a focal search on the high level as well as another focal search planner on the low level that minimizes path conflicts with other agents and therefore decreases the CT size. ECBS specifically mentions how modifying the low level planner to use a weighted cost-to-go heuristic returns paths with many conflicts, leading to a larger CT tree and proved "ineffective in [their] experiments" as direct motivation for reducing the path conflits instead. Explicit Estimation CBS (EECBS) (Li, Ruml, and Koenig 2021) replaces ECBS's high level focal search with Explicit Estimation Search (Thayer and Ruml 2011) but keeps the same low level focal search. To the authors' best knowledge, *no prior work has effectively used a weighted cost-to-go heuristic in any manner in the CBS framework, with the prevailing norm that doing so would lead to more conflicts and reduce performance*.

Our key insight is that we can use the conflict heuristic *along with a weighted cost-to-go heuristic*. We introduce the first bounded suboptimal CBS methods that effectively incorporates a weighted cost-to-go heuristic with the conflict heuristic within CBS's single agent planner. Our contributions are

1. Incorporating the weighted cost-to-go heuristic in the open queue.[1]
2. Effectively weighting cost-to-go heuristic alongside the conflict heuristic in the focal queue, and discovering an important relationship[2] between the two.
3. Proving that PP is a sub-step of suboptimal CBS and showing that weighted suboptimal CBS is the natural generalization[2].

Incorporating Weighted Cost-to-go Heuristic Alongside Conflict Heuristic

CBS utilizes an optimal space-time A* low level planner with a cost-to-go heuristic that measures the optimal distance to goal ignoring conflicts. Bounded sub-optimal CBS methods (e.g. ECBS, EECBS) modify the single agent planner to a focal search that computes a w_{so} sub-optimal path

[1]Additional analysis of how the weights interact with lower bounds and certain CBS improvements is omitted due to space

[2] Supportive experimental results omitted due to space

that minimizes the number of conflicts with other agents (which reduces future constraints in the CT). The low level focal search has two queues; OPEN which searches over optimal paths (sorted by cost) and maintains an optimality bound, and FOCAL which prioritizes w_{so} sub-optimal paths with fewer conflicts (sorted by conflicts). Ties in FOCAL are broken by f_{open}. The user's sub-optimality hyper-parameter w_{so} is assumed to be fixed and outside our optimization. We build upon EECBS as it was shown to outperform ECBS and other MAPF planners, but note that our method is directly usable in ECBS[2] and any other bounded sub-optimal CBS planner using a low level focal planner.

Weighted Open Variant (WO-EECBS)

OPEN's priority function is weighted by w_h, while FOCAL remains unchanged, prioritized by the number of conflicts. To maintain our overall suboptimality bound, the focal bound w_f is scaled to w_{so}/w_h which constrains $w_h \in [1, w_{so}]$ as we need $w_f \geq 1$.

Weighted Focal Variant (WF-EECBS)

We keep OPEN unweighted and instead incorporate the weighted heuristic in FOCAL along with the inadmissible conflict heuristic. This requires us to balance the importance of these competing heuristics via FOCAL's priority function $g + w_h * h + w_c * c$ with $w_h \geq 1, w_c \geq 0$. Changing the magnitude of w_h, w_c changes the relative importance of finding a solution fast (higher w_h) vs avoiding conflicts (higher w_c). Our experimental results[2] discovered an important relationship; the performance is dominated by the ratio $r = w_c/w_h$ rather than the actual w_c or w_h weights. *To highlight the importance of* r*, we reparameterize WF-EECBS in respect to* r *and* w_h *with* $f_{focal}(g, h, c) = g + w_h * (h + r * c)$. Note that $w_h = 1$ and $r \rightarrow \infty$ results in regular EECBS (preferring paths with lowest conflicts). Due to the use of FOCAL, w_h can be arbitrarily large and is not bounded by w_{so}. In our experiments we see that WF-EECBS outperforms WO-EECBS and EECBS, therefore Weighted EECBS (W-EECBS) refers to this weighted focal version.

Lemma 1. *WO-EECBS and WF-EECBS are both* w_{so} *sub-optimal.*

Relating CBS, Prioritized Planning, and W-EECBS

CBS-based algorithms and PP are usually treated as distinct categories of MAPF search based methods. Ma et al. (2019) introduces priorities in CBS as a distinction to regular CBS and does not try to relate the two.

Here we prove that PP is actually equivalent to generating the initial agent paths in the root CT node in EECBS (and other bounded sub-optimal CBS planners like ECBS) with an infinite sub-optimality. With $w_{so} = \infty$ in EECBS, all states in OPEN in the single agent planner are inserted into FOCAL, and therefore expansions are sorted first by their number of conflicts, and then the path f-value. In the root CT node, agents will try to avoid all previous agents and search over all conflict=0 paths, then conflict=1, then conflict=2, etc. This first step is identical to PP; EECBS with $w_{so} = \infty$ differs only in its ability to continue planning over conflicts while PP fails in that scenario. To the authors' knowledge, this is the first time there has been an explicit reduction between sub-optimal CBS and PP. WO-EECBS and WF-EECBS are the two generalized methods combining the weighted low-level planner commonly used in PP with EECBS's conflict resolution mechanism.

Method		Speed up		% faster than Baseline	# Solved
r	w_h	Max	Median		
-	1	1	1	N/A, is baseline	98
-	1.5	37	2.7	70%	101
-	2	32	0.66	39%	67
2.5	8	84	5.5	77%	106
5	4	**131**	9.5	**92%**	**113**
5	16	113	**11**	88%	109
10	8	91	7.2	89%	110

Table 1: We compare the EECBS baseline (first row) with different WO-EECBS w_h parameters (next 2 rows) and different WF-EECBS r, w_h parameters (last 4 rows). We see that WF-EECBS greatly outperforms WO-EECBS and the baseline in the majority of instances.

Experimental Results

We test our methods with different numbers of agents, in increments of 50, on 8 diverse maps from Stern et al. (2019) and report mean values across 5 seeds. We use $w_{so} = 2$ and a timeout of 300 seconds and report the speed up $S_{method} = T_{baseline}/T_{method}$ (larger is better) to normalize differences in hardware, where the baseline is EECBS.

We see that WF-EECBS consistently outperforms WO-EECBS. For WO-EECBS, the effect of w_h on performance fits expectations; larger helps to a certain extent and then hurts due to the interplay with FOCAL. Concretely, WO-EECBS with a "saturated" $w_h = w_{so} = 2$ provides lower speed-ups as FOCAL in that instance has $w_f = w_{so}/w_h = 1$ and has no flexibility to reduce the number of conflicts.

WF-EECBS's performance is most closely related to r as opposed to w_h or w_c. The ratio r explicitly dictates the tradeoff between planning longer to avoid a future conflict or planning shorter and incurring the conflict which will need to be resolved by the constraint tree afterwards. Regular EECBS lacks this flexibility and with $r \rightarrow \infty$ will always plan longer to avoid conflicts. We see WF-EECBS produces large speed-ups and solves more instances than the baseline.

Conclusion

We show two methods of employing a weighted cost-to-go heuristic in suboptimal CBS alongside the conflict heuristic. Our method of effectively incorporating the weighted cost-go-heuristic in FOCAL can produce large speedups at no additional overhead and can be used in other suboptimal CBS planners. We prove that PP is one specific step in suboptimal CBS with infinite suboptimality, and show W-EECBS is the natural generalization of weighted PP and EECBS. **Acknowledgement.** This material is partially supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE1745016 and DGE2140739.

References

Barer, M.; Sharon, G.; Stern, R.; and Felner, A. 2014. Suboptimal variants of the conflict-based search algorithm for the multi-agent pathfinding problem. In *Seventh Annual Symposium on Combinatorial Search*.

Biswas, J.; Mendoza, J. P.; Zhu, D.; Choi, B.; Klee, S.; and Veloso, M. 2014. Opponent-driven planning and execution for pass, attack, and defense in a multi-robot soccer team. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, 493–500.

Boyarski, E.; Felner, A.; Le Bodic, P.; Harabor, D. D.; Stuckey, P. J.; and Koenig, S. 2021. f-Aware Conflict Prioritization amp; Improved Heuristics For Conflict-Based Search. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14): 12241–12248.

Boyarski, E.; Felner, A.; Stern, R.; Sharon, G.; Tolpin, D.; Betzalel, O.; and Shimony, E. 2015. ICBS: Improved conflict-based search algorithm for multi-agent pathfinding. In *Twenty-Fourth International Joint Conference on Artificial Intelligence*.

Čáp, M.; Novák, P.; Kleiner, A.; and Seleckỳ, M. 2015. Prioritized planning algorithms for trajectory coordination of multiple mobile robots. *IEEE transactions on automation science and engineering*, 12(3): 835–849.

Erdmann, M.; and Lozano-Perez, T. 1987. On multiple moving objects. *Algorithmica*, 2(1): 477–521.

Kusnur, T.; Mukherjee, S.; Saxena, D. M.; Fukami, T.; Koyama, T.; Salzman, O.; and Likhachev, M. 2021. A planning framework for persistent, multi-UAV coverage with global deconfliction. In *Field and Service Robotics*, 459–474. Springer.

Li, J.; Felner, A.; Boyarski, E.; Ma, H.; and Koenig, S. 2019. Improved Heuristics for Multi-Agent Path Finding with Conflict-Based Search. In *IJCAI*, volume 2019, 442–449.

Li, J.; Gange, G.; Harabor, D.; Stuckey, P. J.; Ma, H.; and Koenig, S. 2020a. New techniques for pairwise symmetry breaking in multi-agent path finding. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 30, 193–201.

Li, J.; Harabor, D.; Stuckey, P. J.; and Koenig, S. 2021. Pairwise Symmetry Reasoning for Multi-Agent Path Finding Search. *CoRR*, abs/2103.07116.

Li, J.; Ruml, W.; and Koenig, S. 2021. Eecbs: A bounded-suboptimal search for multi-agent path finding. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 12353–12362.

Li, J.; Tinka, A.; Kiesel, S.; Durham, J. W.; Kumar, T. S.; and Koenig, S. 2020b. Lifelong Multi-Agent Path Finding in Large-Scale Warehouses. In *AAMAS*, 1898–1900.

Ma, H.; Harabor, D.; Stuckey, P. J.; Li, J.; and Koenig, S. 2019. Searching with consistent prioritization for multi-agent path finding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, 7643–7650.

Stern, R.; Sturtevant, N. R.; Felner, A.; Koenig, S.; Ma, H.; Walker, T. T.; Li, J.; Atzmon, D.; Cohen, L.; Kumar, T. K. S.; Boyarski, E.; and Bartak, R. 2019. Multi-Agent Pathfinding: Definitions, Variants, and Benchmarks. *Symposium on Combinatorial Search (SoCS)*, 151–158.

Thayer, J. T.; and Ruml, W. 2011. Bounded suboptimal search: A direct approach using inadmissible estimates. In *Twenty-Second International Joint Conference on Artificial Intelligence*.

Velagapudi, P.; Sycara, K.; and Scerri, P. 2010. Decentralized prioritized planning in large multirobot teams. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 4603–4609. IEEE.