

ECBS with Flex Distribution for Bounded-Suboptimal Multi-Agent Path Finding

Shao-Hung Chan,¹ Jiaoyang Li,¹ Graeme Gange,² Peter J. Stuckey,²
Daniel Harabor,² Sven Koenig,¹

¹University of Southern California, USA

²Monash University, Australia

Abstract

1 Multi-Agent Path Finding (MAPF) is the problem of finding
2 collision-free paths for multiple agents on a map. Conflict-
3 Based Search (CBS) is a leading MAPF solver that solves the
4 MAPF instances completely and optimally. Based on CBS,
5 Enhanced Conflict-Based Search (ECBS) is a bounded sub-
6 optimal MAPF solver that uses focal search to speed up CBS.
7 The costs of solutions generated by ECBS are within a given
8 factor away from the optimal. In this paper, we propose Flex-
9 ible ECBS (FECBS), which is an ECBS variant that uses the
10 range between the costs of paths and its upper bound to speed
11 up while still providing solutions within the upper bound. The
12 empirical evaluations show that using our approach can lead
13 to 7.2% higher success rates than the state-of-the-art ECBS
14 variant in the small map, and 10.4% higher in the large map.

1 Introduction

15
16 Multi-Agent Path Finding (MAPF) is the problem of find-
17 ing collision-free paths as a *solution* for multiple agents on
18 the map. The travel time of an agent is known as the *path*
19 *cost* of the agent, and the sum of the travel time is known as
20 the *sum of cost (SoC)*. If a solution whose SoC is minimum,
21 then it is known as an *optimal* solution. On the other hand, if
22 a solution whose SoC is away from the optimal value within
23 a user-specified bound, then it is known as a *bounded sub-*
24 *optimal* solution. Real-world applications of MAPF include
25 autonomous warehouse (Li et al. 2020), UAV traffic man-
26 agement (Ho et al. 2019), video game characters (Ma et al.
27 2017), and office robots (Veloso et al. 2015).

28 Conflict-Based Search (CBS) (Sharon et al. 2015) is a
29 leading two-level MAPF solver for solving MAPF opti-
30 mally. On the low level, it plans paths for single agents.
31 On the high level, it resolves collisions using the best-first
32 search in the space of collision resolutions, with each branch
33 being a new candidate plan that forces one agent or the other
34 find new path to avoid the collision. However, since solv-
35 ing MAPF optimally is a NP-hard problem (Yu and LaValle
36 2013), CBS suffers from the scalability problem, which moti-
37 vates finding bounded-suboptimal solvers to speed up the
38 search. Enhanced CBS (ECBS) (Barer et al. 2014) is a vari-
39 ant of CBS that aims to find the bounded-suboptimal MAPF
40 solution. That is, given user-specified suboptimal factor,

41 ECBS finds the solution with cost at most the factor away
42 from optimal. ECBS replaces the best-first search with *focal*
43 *search* (Pearl and Kim 1982) on both high and low levels of
44 CBS. Once the factor is sufficiently large, leading to suffi-
45 ciently large upper bound, ECBS can quickly find collision-
46 free paths where agents are allowed to take some delays or
47 detours, since it does not have to prove that solutions of
48 lower costs do not exist.

49 While solving MAPF, ECBS may generate candidate
50 plans whose SoC is less than the acceptable upper bound.
51 Thus, in this paper, we proposed Flexible ECBS (FECBS),
52 an ECBS variant that uses the range between the upper
53 bound and SoC to improve the efficiency.

2 Problem Definition

54
55 A MAPF problem, defined in (Stern et al. 2019), consists
56 of an undirected graph $G = (V, E)$ and a set of k agents
57 $\{a_1, a_2, \dots, a_k\}$. Each agent a_i has a unique start vertex
58 $s_i \in V$ and a unique goal vertex $g_i \in V$. Time is discretized
59 into timesteps. At every timestep, an agent performs one of
60 the two actions: the *move* action is to move from the current
61 vertex to an adjacent vertex, and the *wait* action is to stay
62 at the current vertex. An agent continues to exist after it has
63 reached its goal vertex.

64 A path p_i is a sequence of vertices that starts at start vertex
65 s_i and ends at goal vertex g_i . Each vertex in p_i corresponds
66 to the vertex where agent a_i stays at each timestep. Every
67 two adjacent vertices in path p_i are either adjacent vertices
68 in G , meaning that agent a_i *moves*, or identical, meaning
69 that agent a_i *waits*. The path cost of path p_i is the number
70 of timestep, or the traveling time, for agent a_i moving from
71 start vertex s_i to goal vertex g_i , which is also the path length.
72 The timesteps that agent a_i terminally waits at goal vertex g_i
73 are excluded while calculating the cost of path p_i . A *colli-*
74 *sion* (or ,equivalently, *conflict*) between two agents a_i and
75 a_j belongs to one of the categories. It can be a *vertex con-*
76 *flict* $\langle a_i, a_j, v, t \rangle$, where agents a_i and a_j occupy the same
77 vertex $v \in V$ at the same timestep t . It can also be an *edge*
78 *conflict* $\langle a_i, a_j, u, v, t \rangle$, where agents a_i and a_j traverse the
79 same edge $(u, v) \in E$ in opposite directions from timestep
80 t to timestep $t + 1$. A solution is a set of conflict-free paths,
81 one for each agent. An optimal solution is a solution with
82 minimum SoC of the paths. In this paper, the graphs are 4-
83 neighbor grids with vertices corresponding to the unblocked

84 cells and edges corresponding to the connections between
 85 adjacent unblocked cells in the four main compass direc-
 86 tions.

87 3 Existing MAPF Solvers

88 In this section, we introduce Conflict-Based Search (CBS)
 89 and Enhanced CBS (ECBS). Both of them rely on a two-
 90 level architecture where the low level plans paths for sin-
 91 gle agents and the high level resolves conflicts between two
 92 agents.

93 3.1 Conflict-Based Search

94 Before introducing CBS (Sharon et al. 2015), we define *con-*
 95 *straints*, that are used in CBS to resolve conflicts between
 96 two agents. A vertex conflict $\langle a_i, a_j, v, t \rangle$ can be resolved
 97 by prohibiting either agent a_i or agent a_j from staying at
 98 vertex v at timestep t , leading to *vertex constraints* $\langle a_i, v, t \rangle$
 99 and $\langle a_j, v, t \rangle$. An edge conflict $\langle a_i, a_j, u, v, t \rangle$, on the other
 100 hand, can be resolved by prohibiting either agent a_i or agent
 101 a_j from traversing edge (u, v) in the opposite direction from
 102 timestep t to timestep $t + 1$, leading to *edge constraints*
 103 $\langle a_i, u, v, t \rangle$ and $\langle a_j, v, u, t \rangle$.

104 On the low level, CBS plans the path of agent a_i by view-
 105 ing vertex $v \in V$ and timestep t as a spatio-temporal node
 106 $n = (v, t)$ and runs A* to find a minimum cost path that
 107 satisfies the constraints given by the high level. That is, A*
 108 on the low level uses an open list OPEN and sorts the spatio-
 109 temporal nodes in OPEN in increasing order of their f val-
 110 ues $f^i(n) = g^i(n) + h^i(n)$, where $g^i(n)$ is the number of
 111 timesteps for agent a_i to move from its start vertex s_i to ver-
 112 tex v and $h^i(n)$ is an admissible heuristic that estimates the
 113 distance from vertex v to its goal vertex g_i . The A* on the
 114 low level breaks ties in favor of a path that has the fewest
 115 conflicts with the paths of other agents.

On the high level, CBS runs a best-first search on a binary
constraint tree (CT). Each CT node N has two components:
 (1) a set of paths of all agents generated by the low-level
 search $N.paths$, with the path of agent a_i being $N.paths[i]$
 and its cost being $|N.paths[i]|$, and (2) a set of vertex and
 edge constraints, $N.constraints$, that coordinate agents to
 avoid conflicts. The *cost* of CT node N , $N.cost$, is the SoC
 of $N.paths$, that is,

$$N.cost = \sum_{i=1}^k |N.paths[i]|. \quad (1)$$

116 CBS begins its high-level search at the *root* CT node, that
 117 contains a minimum cost path for each agent and an empty
 118 set of constraints. While expanding CT node N , if there are
 119 no conflicts among the paths of N , CBS returns the paths
 120 of N and terminates. Otherwise, it picks one of the conflicts
 121 and resolves it by *branching*, also known as splitting N into
 122 two child CT nodes. In each child CT node, CBS adds a
 123 vertex or edge constraint to one of the conflicting agents to
 124 prohibit it from utilizing the conflicted vertex or edge, re-
 125 spectively, at the conflicted timestep. It then performs A*
 126 on the low level to replan the path of the agent with the
 127 additional constraint and leaves all other paths unchanged.
 128 In other words, CBS only replans one agent in a child CT

node. CBS solves MAPF optimally by performing best-first
 searches on both high and low levels.

514 3.2 Enhanced CBS

Enhanced CBS (ECBS) (Barer et al. 2014) uses focal search
 on both the high and low levels to speed up CBS signifi-
 cantly, which is a bounded suboptimal search instead of the
 best-first search. On the low level, ECBS uses an open list
 OPEN and sorts its spatio-temporal nodes n in increasing
 order of their f values $f_1^i(n)$, which is identical to func-
 tion $f^i(n)$ of CBS, to find a path for agent a_i . Let $best^i$
 be the node n with the minimum $f_1^i(n)$ value in OPEN
 and w be the user-specified suboptimality factor. At the be-
 ginning of the low-level focal search, only the *root spatio-*
temporal node $n_i^0 = (s_i, 0)$ is in OPEN, thus $n_i^0 = best^i$
 and the *initial minimum* $f_1^i(n)$ value is set to $f_1^i(n_i^0)$. The
 low-level focal search also uses a focal list FOCAL that con-
 tains all spatio-temporal nodes $n = (v, t)$ in OPEN with
 $f_1^i(n) \leq w \cdot f_1^i(best^i)$. The spatio-temporal nodes in FO-
 CAL are sorted in increasing order of their different f val-
 ues $f_2^i(n)$, which specifies the number of conflicts of the
 path of agent a_i with the paths of other agents in $N.paths$
 while agent a_i moves from its start vertex s_i to vertex v . The
 low-level focal search expands a node n with the minimum
 $f_2^i(n)$ value in FOCAL. We define the *low-level lower bound*
 on the cost of the optimal path of agent a_i as $N.lb[i]$. Since
 $f_1^i(n)$ uses an admissible heuristic, $N.lb[i] = f_1^i(best^i)$ for
 ECBS. Thus, the low-level focal search always returns a path
 for agent a_i with a cost of at most w times the optimal path
 cost c_i^* , meaning that,

$$N.lb[i] \leq |N.paths[i]| \leq w \cdot N.lb[i] \leq w \cdot c_i^*. \quad (2)$$

Here, we let $N.ub[i] = w \cdot f_1^i(best^i)$ as the *low-level upper*
bound of agent a_i in CT node N . The low-level focal search
 also returns a lower bound on the cost of the optimal path
 for agent a_i , which is the $f_1^i(best^i)$ value when the low-level
 search terminates.

On the high level, ECBS also runs focal search on CT.
 Given a CT node N , we define its lower-bound value as
 $LB(N) = \sum_{i=1}^k f_1^i(best^i)$, and we define its upper-bound
 value as $UB(N) = w \cdot \sum_{i=1}^k f_1^i(best^i)$. We define the num-
 ber of conflicts between all the pairs of paths in CT node
 N to be $f_2(N)$. Let $LB = \min(LB(N) \mid N \in OPEN)$.
 Since $LB(N)$ is a lower bound on the minimum SoC of
 the solutions below CT node N , LB is a lower bound on
 the optimal cost, denoted as C^* . The high-level focal search
 sorts the CT nodes N in OPEN according to $LB(N)$ and
 adds the CT nodes N with the costs of at most $w \cdot LB$ into
 FOCAL, where w is the same suboptimality factor as the
 low level. ECBS expands the CT node with the minimum f_2
 value in FOCAL, namely the CT node with the fewest num-
 ber of conflicts. Given a MAPF instance with optimal cost
 C^* , since ECBS only selects CT nodes whose costs are at
 most $w \cdot LB$ and $LB \leq C^*$, it finds a solution whose cost is
 at most $w \cdot C^*$, that is,

$$LB \leq N.cost \leq w \cdot LB \leq w \cdot C^*. \quad (3)$$

Here, we define $UB = w \cdot LB$ as the upper bound of high-
 level focal search. We define UB as the high-level upper

bound. As the value of w increases, ECBS has more sub-optimal solutions to choose from, which might decrease its runtime.

4 Flexible ECBS (FECBS)

In ECBS, each CT node N has path $N.paths[i]$ of agent a_i and the low-level lower bound of the path $N.lb[i]$. Given a user-specified factor w , the low-level upper bound of path $N.paths[i]$ is $w \cdot N.lb[i]$. We define the *flex* of agent a_i , or $N.flex[i]$, as the difference between the low-level upper bound of the path and the path cost, that is,

$$N.flex[i] = w \cdot N.lb[i] - |N.paths[i]|. \quad (4)$$

One of the limitations of ECBS is that each agent a_i , while being replanned on the low level, only considers its upper bound given by its own lower bound $f_1^i(best^i)$. Thus, we propose Flexible ECBS (FECBS), an ECBS variant that increases the low-level upper bound of agent a_i by adding the flex values from other agents, known as the *flex distribution* scheme. The low-level upper bound of agent a_i of FECBS is the summation of $w \cdot N.lb[i]$ and the sum of the flex values from other agents, that is,

$$N.ub[i] = w \cdot N.lb[i] + \sum_{a_j \in A \setminus a_i} N.flex[j], \quad (5)$$

where $\sum_{a_j \in A \setminus a_i} N.flex[j]$ is the sum of the flex values from other agents aside from agent a_i , indicating we distribute all the flex values from other agents to the agent that is going to be replanned. Thus, the low-level focal search of FECBS adds spatio-temporal node n whose $f_1^i(n)$ is within low-level upper bound $N.ub[i]$ to FOCAL. Also, FECBS should keep track of the maximum low-level lower bound of agents along the branch of CT during the search. Supposed that the parent CT node of CT node N is \hat{N} and the low-level lower bound of agent a_i in CT node \hat{N} is $\hat{N}.lb[i]$. While replanning agent a_i in CT node N , if the sum of the flex values from other agents is large enough, there may be a case where $f_1^i(best^i)$ is less than the low-level lower bound of agent a_i in CT node \hat{N} , while the low-level upper bound of agent a_i in CT node N is greater than the low-level upper bound of agent a_i in its parent CT node \hat{N} , that is,

$$f_1^i(best^i) < \hat{N}.lb[i] \quad (6)$$

$$w \cdot f_1^i(best^i) - |N.paths[i]| > w \cdot \hat{N}.lb[i] - |\hat{N}.paths[i]| \quad (7)$$

In this case, the low-level focal search of FECBS may find a path of agent a_i that satisfies the constraints while $f_1^i(best^i)$ is less than the lower bound of agent a_i in the parent CT node \hat{N} . Thus, if we set $f_1^i(best^i)$ to $N.lb[i]$ directly as what ECBS does, the lower-bound value of CT node N is less than the lower-bound value of its parent CT node \hat{N} , meaning that the sum of the low-level lower bounds of agents decreases. This results in decreasing of $w \cdot N.lb[i]$, and flex that have been used to replan any agents before branching to CT node N may “disappear”, which causes FECBS to be invalid. Thus, given a branch of CT B and N being its tail, we define the maximum low-level lower bound of agents along




Figure 1: Maps used in the experiments with their size given in the form $height \times width$.

branch B as $B.lb[i]$, and the initial minimum $f_1^i(n)$ value of agent a_i in CT node N should be

$$\max(f_1^i(n_i^0), B.lb[i]). \quad (8)$$

While replanning agent a_i on the low-level, the sum of flex values from other agents may be negative, meaning that there are other agents that have been replanned previously by using flex value of agent a_i . Still, the low-level upper bound of agent a_i should be at least the low-level lower bound of agent a_i , meaning that,

$$N.lb[i] \leq N.ub[i]. \quad (9)$$

If $N.lb[i] = N.ub[i]$, then this means that other agents have use up all the flex value of agent a_i . Thus, FECBS will run the low-level focal search exactly the same as the low-level focal search of ECBS with the initial minimum value of agent a_i in CT node N being as Equation (8). FECBS potentially distributes the flex values to agents whose constraints are difficult to satisfy. Also, for agent whose flex has been used up for other agents, FECBS tends to increase its low-level lower bound to find a path that satisfies the constraints, resulting in the increment of both high-level and low-level upper bounds and even generates more flex for other agents. Since the flex values are obtained from the range between the high-level upper bound and SoC, FECBS still guarantees to be bounded suboptimal with user-specified suboptimal factor w .

5 Empirical Evaluation

As shown in Figure 1, the scenarios for these experiments are all 4-neighbor grids from the MAPF benchmark suite (Stern et al. 2019), including a 32×32 grid map with 20% blocked cells (Grid), and a 257×256 grid map from the video game Dragon Age: Origin (DAO) (Den520d). We use both the “even” and “random” scenarios from the benchmark, with each containing 25 instances for each map and each number of agents. The start and goal vertices are distributed evenly in the instances of the “even” scenario. The start and goal vertices are distributed randomly in the instances of the “random” scenario. Our main comparison metric is the *success rate*, which is the percentage of the instances solved within a runtime limit of 5 minutes. For Grid map, which has small size but high obstacle density, we set

Maps	#Runs	ECBS	FECBS	ECBS (RR)	FECBS (RR)
Grid	5	52.4	63.8	59.2	67.2
	20			59.8	68.2
	30			62.8	70
	40			62.6	69.8
Den520d	5	76.4	86.8	60.4	58.4
	20			42	48.4
	30			54	57.2
	40			62	76.8

Table 1: The success rate (in percentage) of ECBS variants with $\#Runs \in \{5, 20, 30, 40\}$.




Figure 2: Success rate (%) of ECBS, ECBS (RR) with its best #Runs, FECBS, and FECBS (RR) with its best #Runs on different maps.




Figure 3: Success rate (%) of ECBS, ECBS (RR) with its best #Runs, FECBS, and FECBS (RR) with its best #Runs on different maps.

	ECBS	FECBS	ECBS (RR)	FECBS (RR)
Grid	1.027	1.028	1.031	1.029
Den520d	1.006	1.0062	1.0057	1.0058

Table 2: The average suboptimalities of MAPF solver in different maps.

its weakness running in the large map and large number of agents. Still, our FECBS shows advantage among all the ECBS variants, which is also shown in Table 1.

We also analyze the generated CT nodes of ECBS variants to show the efficiency while solving MAPF instances. If a MAPF solver is able to solve a MAPF instance with fewer CT nodes than another MAPF solver, then it has higher efficiency. In Grid map, we compare FECBS (RR) and ECBS (RR) with their individual best #Runs. In Den520d map, we compare FECBS and ECBS, since the rapid random restart scheme does not show improvements. Figure 3 shows the comparison of generated CT node between two ECBS variants. We use a logarithmic scale for both axes. For instances that not solved within the runtime limit, we set the number of CT nodes generated by the MAPF solver to ∞ . Instances on the right hand side of the dashed line are the ones that can be solved by FECBS (RR) (or FECBS) with fewer generated CT nodes than by ECBS (RR) (or ECBS). Figure 3 shows that ECBS variants with the flex distribution scheme can reduce the number of generated CT nodes and thus improve the efficiency.

To analyze how close between the SoC of the solution and the high-level upper bound, we use the *suboptimality* α , which is the ratio between the SoC and the high-level lower bound of a MAPF instance. Thus, if the suboptimality is close to the user-specified suboptimality factor w , then the SoC of the solution is close its high-level upper bound. We also define *average suboptimality* $\bar{\alpha}$ as the average of the suboptimalities over the MAPF instances been solved. Table 2 shows the average suboptimality of the ECBS variants over the instances in different maps. Although flex distribution seems to push the SoC to the high-level upper bound, the average suboptimalities of FECBS and FECBS (RR) are still close to ECBS and ECBS (RR). Thus, the flex distribution scheme does not affect the solution qualities much.

the suboptimality bound to $w = 1.05$. For Den520d map, $w = 1.05$ is too large, resulting in all ECBS variants having high success rates and preventing us from distinguishing among them. Thus, we set $w = 1.01$. We implement our ECBS variants in C++ and run them on servers with 2.80 GHz Intel Xeon Processors E5-2640 v4 and 2 GB RAM.

We compare our FECBS with Rapid Randomized Restart ECBS (Cohen et al. 2018), known as ECBS (RR), which is the state-of-the-art ECBS variant. ECBS (RR) uses a user-specified *number of runs* (or #Runs) to determine the times for restarting the search. Once the time limit T and #Runs are set, ECBS (RR) will restart every $T/\#Runs$, each time randomly shuffling the order of agents before the restart. We evaluate ECBS (RR) on the maps with #Runs $\in \{5, 20, 30, 40\}$. We define the *best #Runs* as the value of #Runs in the set that leads to the highest number of solved instances. Since rapid randomized restart scheme only focuses on determining when to restart the search, it can also be used in FECBS, resulting in FECBS (RR). Table 1 shows the number of solved instances for different values of #Runs.

Figure 2 shows success rate versus the number of agents (or #Agents) of ECBS, ECBS (RR), FECBS, and FECBS (RR). We only show ECBS (RR) and FECBS (RR) with their individual best #Runs. In Grid map, as the number of agent increases, FECBS has higher success rate than ECBS and is compatible with ECBS (RR) of the best #Runs. Also, the combination of the flex distribution and the rapid random restart schemes even improves the success rates. In Den520d map, if #Runs is too high, then ECBS (RR) shows

6 Conclusions

238 In this paper, we introduce a new scheme for ECBS to
 239 use the range between its high-level upper bound and SoC,
 240 called the flex distribution scheme, and propose FECBS and
 241 FECBS (RR). Both ECBS variants gather flex values from
 242 other agents to increase the low-level upper bound of the
 243 agent that is currently being replanned. Potentially, con-
 244 straints of the agents may be satisfied with the increment of
 245 the low-level lower bounds. The empirical evaluations show
 246 that using our approach can lead to 7.2% higher success
 247 rates than the state-of-the-art ECBS variant (ECBS (RR)) in
 248 Grid map, and 10.4% higher in Den520d map. The future
 249 work for this paper may be designing heuristics for flex dis-
 250 tribution, instead of assigning all the flex values to a single
 251 agent, to improve the efficiency.

References

- 252 Barer, M.; Sharon, G.; Stern, R.; and Felner, A. 2014. Subopti-
 253 mal Variants of the Conflict-Based Search Algorithm for the Multi-
 254 Agent Pathfinding Problem. In *Proceedings of the Annual Symposi-
 255 um on Combinatorial Search (SoCS)*, 19–27. 256
- 257 Cohen, L.; Koenig, S.; Kumar, T. K. S.; Wagner, G.; Choset, H.;
 258 Chan, D.; and Sturtevant, N. 2018. Rapid Randomized Restarts for
 259 Multi-Agent Path Finding: Preliminary Results. In *Proceedings
 260 of the 17th International Conference on Autonomous Agents and
 261 MultiAgent Systems (AAMAS)*, 1909–1911. 261
- 262 Ho, F.; Salta, A.; Geraldles, R.; Goncalves, A.; Cavazza, M.; and
 263 Prendinger, H. 2019. Multi-Agent Path Finding for UAV Traffic
 264 Management. In *Proceedings of the International Conference on
 265 Autonomous Agents and MultiAgent Systems (AAMAS)*, 131–139. 265
- 266 Li, J.; Tinka, A.; Kiesel, S.; Durham, J. W.; Kumar, T. K. S.; and
 267 Koenig, S. 2020. Lifelong Multi-Agent Path Finding in Large-
 268 Scale Warehouses. In *Proceedings of the International Confer-
 269 ence on Autonomous Agents and MultiAgent Systems (AAMAS)*,
 270 1898–1900. 270
- 271 Ma, H.; Yang, J.; Cohen, L.; Kumar, T. K. S.; and Koenig, S. 2017.
 272 Feasibility Study: Moving Non-Homogeneous Teams in Congested
 273 Video Game Environments. In *Proceedings of the 13th AAAI Con-
 274 ference on Artificial Intelligence and Interactive Digital Entertain-
 275 ment (AIIDE)*, 270–272. 275
- 276 Pearl, J.; and Kim, J. H. 1982. Studies in Semi-Admissible Heuris-
 277 tics. *IEEE Transactions on Pattern Analysis and Machine Intelli-
 278 gence* 4(4): 392–399. 278
- 279 Sharon, G.; Stern, R.; Felner, A.; and Sturtevant, N. R. 2015.
 280 Conflict-Based Search for Optimal Multi-Agent Pathfinding. *Arti-
 281 ficial Intelligence* 219: 40–66. 281
- 282 Stern, R.; Sturtevant, N. R.; Felner, A.; Koenig, S.; Ma, H.; Walker,
 283 T.; Li, J.; Atzmon, D.; Cohen, L.; Kumar, T. K. S.; Boyarski, E.; and
 284 Bartak, R. 2019. Multi-Agent Pathfinding: Definitions, Variants,
 285 and Benchmarks. In *Proceedings of the International Symposium
 286 on Combinatorial Search (SoCS)*, 151–159. 286
- 287 Veloso, M. M.; Biswas, J.; Coltin, B.; and Rosenthal, S. 2015.
 288 CoBots: Robust Symbiotic Autonomous Mobile Service Robots.
 289 In *Proceedings of the 24th International Joint Conference on Arti-
 290 ficial Intelligence (IJCAI)*, 4423–4429. 290
- 291 Yu, J.; and LaValle, S. M. 2013. Structure and Intractability of
 292 Optimal Multi-Robot Path Planning on Graphs. In *Proceedings of
 293 the AAAI Conference on Artificial Intelligence*, 1443–1449. 293