

Optimizing Crowd-Aware Multi-Agent Path Finding through Local Communication with Graph Neural Networks

Phu Pham and Aniket Bera

Department of Computer Science, Purdue University, USA

Project page: <https://ideas.cs.purdue.edu/research/projects/cramp>

Abstract—Multi-Agent Path Finding (MAPF) in crowded environments presents a challenging problem in motion planning, aiming to find collision-free paths for all agents in the system. MAPF finds a wide range of applications in various domains, including aerial swarms, autonomous warehouse robotics, and self-driving vehicles. Current approaches to MAPF generally fall into two main categories: centralized and decentralized planning. Centralized planning suffers from the curse of dimensionality when the number of agents or states increases and thus does not scale well in large and complex environments. On the other hand, decentralized planning enables agents to engage in real-time path planning within a partially observable environment, demonstrating implicit coordination. However, they suffer from slow convergence and performance degradation in dense environments. In this paper, we introduce CRAMP, a novel crowd-aware decentralized reinforcement learning approach to address this problem by enabling efficient local communication among agents via Graph Neural Networks (GNNs), facilitating situational awareness and decision-making capabilities in congested environments. We test CRAMP on simulated environments and demonstrate that our method outperforms the state-of-the-art decentralized methods for MAPF on various metrics. CRAMP improves the solution quality up to 59% measured in makespan and collision count, and up to 35% improvement in success rate in comparison to previous methods.

I. INTRODUCTION

Multi-Agent Path Finding (MAPF) poses challenges with broad applications in autonomous warehouses, robotics, aerial swarms, and self-driving vehicles [1], [2], [3], [4], [5], [6], [7]. The objective is to plan paths for multiple agents to navigate from start to goal positions in obstacle-laden environments. A critical constraint of such systems is to ensure agents navigate concurrently without collisions. Two main categories of approaches are centralized and decentralized planning. Centralized planners [8], [9], [10], [11] aim to find optimal solutions that minimize the length of collision-free paths. They are effective in small and sparse environments but face limitations in real-time performance and scalability in large and dense environments [12], [13]. These methods require complete knowledge of the environment and full replanning when a change occurs, leading to exponential computation times with increased agents, obstacles, and world size. Recent studies [14], [11] have sought to discover real-time solutions. However, these solutions remain sub-optimal and still require access to global information about the world.

On the contrary, decentralized methods [15], [16], [17], [18], [19], [1] seek to tackle these challenges by allowing each agent to acquire their own policies. In these approaches,

The authors are with the Department of Computer Science, Purdue University, USA, {[pham84](mailto:pham84@purdue.edu), [aniketbera](mailto:aniketbera@purdue.edu)}@purdue.edu




Fig. 1. Comparison of paths between our CRAMP approach and the previous state-of-the-art methods. Colored circles with numerical labels represent the starting positions of agents, while dashed circles with corresponding numbers denote the target positions. The obstacles are marked with grey squares. Our innovative crowd-aware method for multi-agent path-finding challenges significantly outperforms existing approaches such as PRIMAL [15] and PICO [17], yielding notably shorter solutions.

agents can reactively plan paths within partially observable environments. Such methods prove beneficial in situations where agents lack comprehensive knowledge of the world, as is often the case in the domain of autonomous vehicles. Rather than pursuing an optimal solution for all agents, decentralized planners train local policies that rapidly generate sub-optimal solutions as a tradeoff between speed and solution quality. Given that agents make their decisions based on local information, decentralized approaches often face challenges in achieving effective global coordination among agents. In cluttered, dynamic environments characterized by congestion or rapid changes, agents may tend to prioritize their individual goals, potentially resulting in conflicts and inefficiencies that affect the overall system's performance.

To tackle the challenges presented by dense crowds, our approach, CRAMP, integrates a crowd-awareness mechanism and local communication via Graph Neural Networks (GNN) [20]. This strategy is guided by a boosted curriculum training approach, which gradually exposes agents to increasingly complex scenarios. CRAMP is specifically designed to train

intelligent agents capable of efficiently navigating in dense and dynamic environments. By formulating the MAPF problem as a sequential decision-making task, we employ deep reinforcement learning techniques, enabling agents to make optimal decisions while considering the presence of other agents and the dynamic nature of the environment.

We evaluate the CRAMP approach through extensive experiments on various synthetic environments, comparing it against state-of-the-art MAPF algorithms. Our results demonstrate that CRAMP achieves superior performance in terms of solution quality and computational efficiency while maintaining robustness in crowded settings. Moreover, we showcase the adaptability of our approach in dynamic and evolving environments. Figure 1 illustrates a comparison of computed paths between PRIMAL [15], PICO [17], and our method, CRAMP. Our approach significantly shortens the path to the target position while mitigating potential congestion by prioritizing actions that lead to less densely populated areas. This strategy effectively reduces the likelihood of deadlocks and enhances overall path efficiency.

Our main contributions can be summarized as follows:

- We present a reinforcement learning approach for the multi-agent path-finding problem, utilizing a carefully designed crowd-aware mechanism that enables agents to adapt their policies dynamically based on crowd density and behavior.
- Exploiting local communication between nearby agents via Graph Neural Networks, we facilitate agents to gather information from and predict the behavior of nearby agents, thereby enhancing their decision-making capabilities.
- We conduct a series of experiments, conclusively showcasing that our approach surpasses state-of-the-art methods across a range of metrics, including success rate, collision count, and makespan.

II. RELATED WORK

Researchers have developed a wide range of algorithms and techniques to solve the MAPF problem, aiming to strike a balance between optimality, speed, and adaptability to varied maps and conditions. Centralized approaches have traditionally been favored for MAPF, but they come with several drawbacks. As the number of agents and the complexity of the environment increase, centralized algorithms may become computationally intractable, leading to long planning times [10], [21]. Furthermore, they are susceptible to single points of failure, where if the central planner fails, the entire system may collapse.

Decentralized planners have garnered attention due to their ability to mitigate the computational complexity associated with centralized planners and their resilience against single points of failure. However, these approaches may face challenges in managing communication overhead and ensuring effective coordination, especially in scenarios with a large number of agents [22], [23], [15].

Moreover, the exploration of hybrid techniques that integrate both centralized and decentralized components remains an important avenue of research [24], [25]. Hybrid approaches could potentially capitalize on the strengths of both paradigms while mitigating their respective limitations.

In this work, we primarily study the centralized and decentralized approaches. Future research could explore hybrid approaches to gain a more comprehensive understanding of the tradeoffs involved and to further advance the field.

A. Centralized approaches

Centralized approaches for MAPF involve using a centralized planner to find optimal or near-optimal paths for all agents simultaneously. These approaches are suitable for scenarios where global coordination among agents is feasible and desired. Among the centralized methods, Conflict-based search (CBS) [10] is a widely adopted algorithm that decomposes the MAPF problem into single-agent path-finding problems while considering agent conflicts. It iteratively identifies conflicts, prioritizes them, and resolves them by searching for alternative paths or time steps for the conflicting agents. Other variants of CBS [26], [27], [28], [29] extend CBS to handle dynamic environments and continuous time.

Extending A* [30] to MAPF, hierarchical A* [31] and M* [32] plan individual paths for each agent and then search forward in time to detect potential collisions. When collisions occur, they perform joint planning through limited backtracking to resolve the collision and continue with individual agent plans. One variant, ODrM* [33], reduces the need for joint planning by breaking down the problem into independent collision sets and employs Operator Decomposition (OD) [34] to manage search complexity effectively.

Another line of approach is sampling-based methods [35], [36], [37], which have also been adapted to MAPF. These methods generate random samples of possible agent configurations and construct paths by connecting these samples while avoiding collisions.

Recently, MAPF-LNS [13] and MAPF-LNS2 [38] solvers can achieve near-optimal solutions in a short amount of time using Large Neighborhood Search. LaCAM [11] proposes a two-level search algorithm to solve large MAPF instances with up to thousands of agents in a matter of seconds. All the above-mentioned approaches require global information.

B. Decentralized approaches

Recent studies have directed their attention towards decentralized policy learning [23], [15], [22], [17], [18], [19], [39], in which individual agents learn their own policies. [23] models the MAPF problem as a Partially Observable Markov Decision Process (POMDP) and uses hierarchical graph-based macro-actions to produce robust solutions. A macro-action is a higher-level action composed of a sequence of atomic actions that an agent can take in its environment.

Additionally, [22] proposes a method called WSCaS (Walk, Stop, Count, and Swap) that optimizes the initial solutions provided by A*. Each agent a_i initially plans a path $\mathcal{P}_i = \{p_0, \dots, p_n\}$ from the start position s_i to the goal position g_i using the A* algorithm. At each position p_j on the path \mathcal{P}_i , the agent a_i utilizes local communication to reactively adjust the path according to the current environment to avoid collisions and deadlocks until it reaches g_i .

In contrast to conventional methods, [40] formulates the MAPF problem within a game theory setting, where each agent tries to maximize their utility function while respecting safety constraints. In this framework, each agent simulates a good learned policy while avoiding the bad ones.




Fig. 2. The network architecture to train the distributed local policies. The network takes local observation of size $10 \times 10 \times 4$ as inputs. The network’s outputs include the predicted probabilities for the actions (size 5×1), predicted return value, validity of predicted action and blocking probabilities (all of size 1×1).

C. Multi-agent reinforcement learning

The trajectory of multi-agent reinforcement learning (MARL) research features key contributions that have systematically addressed the complexities of agent coordination and navigation. PRIMAL [15] focuses on integrating reinforcement learning with imitation learning to enhance local policy development and inter-agent coordination. PICO [17] introduced prioritized communication learning methods, optimizing information exchange and decision-making processes among agents. Additionally, extensions like CPL [18] have adopted a curriculum learning framework, applying distinct loss functions at varying curriculum levels to progressively refine agent policies and tackle the challenges of learning in highly dynamic settings.

D. Graph neural networks for MAPF

Graph Neural Networks (GNNs) have gained significant attention and adoption within multiagent systems due to their demonstrated efficacy in various tasks [41], [42], [43], [44], [45]. In the MAPF settings, GNNs offer a compelling solution by facilitating complex interaction modeling, local communication, and adaptation to dynamic environments. By representing agents and their surroundings as nodes and edges in a graph structure, GNNs enable efficient coordination and decision-making among agents.

Despite the promising potential, limited work exists in applying GNNs to MAPF. Notably, [41] integrated GNNs into MAPF but trained the model through imitation learning, leading to an imbalance between exploration and exploitation. As a consequence, the model’s generalization in new environments is compromised. In this work, we aim to address this issue by integrating GNNs into our framework, seeking to enhance the generalization capabilities of our model. Our contribution to this landscape includes a novel crowd-avoidance mechanism supported by local communication through GNNs. We adopt a simple yet effective curriculum-driven training strategy to enhance learning efficiency and significantly reduce incidents of collisions and deadlocks.

III. OUR APPROACH

This section outlines CRAMP, our approach to addressing the MAPF problem in densely populated environments. The methodology can be divided into several key components, including world modeling, policy learning, a crowd-aware reward function, GNN-based local communication, and a boosted curriculum training strategy.

A. World Modeling

Like most reinforcement learning techniques, we formulate the MAPF problem in a discrete grid-world environment. In this setting, agents partially observe the world within a limited local region of size $m \times m$ centered around themselves. Within this local region, each agent can perceive its own goal, the presence of nearby agents, and the directions to those agents’ goals. These observations are used to train the local policy network.

The action space of each agent includes five discrete actions in the grid world: moving one cell in one of the four directions (N, E, S, W) or remaining stationary. An action is considered valid if it avoids collisions with obstacles and other agents or goes beyond the world’s boundaries.

B. Reinforcement learning for local policy training

We train the agent’s policy using deep reinforcement learning and demonstration learning techniques. Multiple agents are independently trained in a shared environment utilizing the Asynchronous Advantage Actor-Critic (A3C) algorithm [46]. The inputs to the policy network comprise four layers of an agent’s local observation, which include the grid world, obstacles, neighboring agents, and neighbors’ goals, alongside its own goal position. The network’s outputs consist of the next action, predicted return value, validity of the action, and blocking probability. The detailed architecture of the policy network is illustrated in Fig. 2. The CNN feature extractor employs two sequential blocks of [Conv2D-Conv2D-Conv2D-Relu-MaxPool]. The CNN’s output is a 512×1 vector. This feature vector is subsequently input into an LSTM to encode temporal features. In the GNN block, the adjacency matrix is computed based on the local field of view (FOV), where the ego agent is connected to all other agents within the FOV. For each agent within the ego agent’s FOV, encoded feature vectors are extracted to create a node embedding matrix of size $N \times 512$, where N represents the maximum number of visible agents. The GNN processes both the adjacency and node embedding matrices to produce an $N \times 128$ matrix, with each row reflecting a graph signal from each agent. These nodes are averaged and concatenated with the LSTM feature vector. The resulting concatenated feature vector is fed into fully-connected layers to predict the variables.

These outputs are updated after every batch of $T = 256$ steps or after each episode. An episode is finished when all agents successfully arrive at their goals or when the maximum number of episode steps has been reached.

Consider a timestep t ; the network aims to maximize the discounted return for the trajectory starting at t :

$$\mathcal{R}_t = \sum_{i=0}^{\infty} \gamma^i r_{t+i} \quad (1)$$

where γ is the discount factor and r_{t+i} is the reward received by the agent at timestep $t+i$ after taking a series of actions. The optimization can be achieved by minimizing the loss function:

$$\mathcal{L}_{\text{value}} = \sum_{t=0}^T (V(o_t) - \mathcal{R}_t)^2 \quad (2)$$

where $V(o_t)$ is the predicted value of the policy network at timestep t .

The policy is updated by adding the entropy term $H = -\sum_a \pi(a_t = a|o_t) \log(\pi(a_t = a|o_t))$ [47] to the policy loss:

$$\mathcal{L}_\pi = \epsilon H - \sum_{t=0}^T \log(P(a_t|\pi, o)A(o_t, a_t)) \quad (3)$$

in which ϵ is the small weight for the entropy term H , $P(a_t|\pi, o)$ is the probability of taking action a at timestep t with observation o by following the policy π , and $A(o_t, a_t)$ is the approximation of the advantage function.

$$A(o_t, a_t) = \sum_{i=0}^{k-1} \gamma^i r_{t+i} + \gamma^k V(o_{k+t}) - V(o_t) \quad (4)$$

An agent a_i is considered blocking another agent a_j if a_i is on the optimal path of a_j to the goal from its current position and forces a_j to take a detour at least ten steps longer. The optimal paths of these agents are determined using the A^* algorithm. The **blocking** output is used to predict the blocking probability. We use binary cross-entropy loss as the loss function for this term.

$$\mathcal{L}_{\text{block}} = -\sum_{i=1}^T y_i \log(p_i) + (1 - y_i) \log(1 - p_i) \quad (5)$$

where y_i represents the ground truth binary label indicating whether the ego agent is blocking any other agents, p_i is the predicted blocking probability.

The same binary cross-entropy loss is applied to the action validity loss $\mathcal{L}_{\text{valid}}$. The total loss is then formulated as the sum of all losses:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_\pi + \mathcal{L}_{\text{value}} + \mathcal{L}_{\text{blocking}} + \mathcal{L}_{\text{valid}} \quad (6)$$

C. Demonstration learning

Previous studies [48], [49], [50], [15] show that learning from demonstrations can facilitate agents to confine the high-reward regions quickly. The agents then use reinforcement learning to explore these regions to improve their policies. Adopting this approach from PRIMAL [15], we also use ODrM* [33] as the expert agent. The agents will first learn from the ODrM* expert and then randomly switch between exploration and demonstration modes to update their local policies.

D. Crowd-aware reward function

In reinforcement learning for grid-world MAPF, the reward system is simple yet effective: agents lose points for each timestep away from their goal, motivating them to find the fastest route. They are also penalized for idleness before reaching their goal, encouraging exploration and action. Major penalties are given for collisions, exiting the environment, or hitting obstacles, discouraging such actions.

Decentralized policies might result in agents behaving selfishly, blocking others from their goals, and potentially causing deadlocks in tight areas. To foster teamwork, agents receive a significant reward when all complete an episode together but face penalties for hindering others, promoting cooperation and efficiency in the multi-agent system.

Even though these rewards are widely used in MAPF, they often overlook the issue of traffic congestion. None of the previously mentioned reward terms consider agents' behavior in densely populated areas. Typically, deadlocks often happen

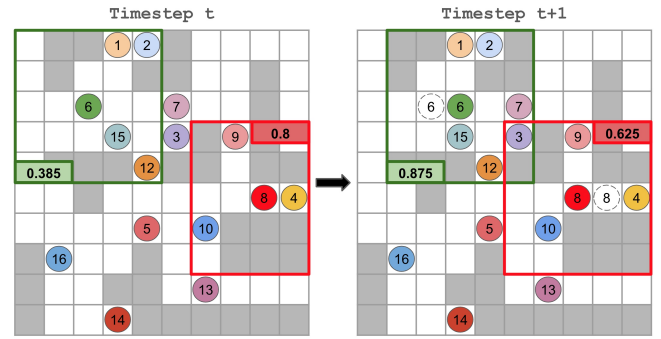


Fig. 3. Example of agents getting negative and positive rewards by moving in or out of a ζ -crowded region. The colored circles with numerical labels are the agents and the dashed circles represent the previous positions.

in congested regions, where the available moving space is severely limited. To reduce the likelihood of a deadlock, we introduce an innovative reward term, denoted as R_n , that could boost agents' performance in terms of deadlock rates. We define an event $\mathcal{E} = \{\delta_a \geq \zeta\}$, where δ_a is the density of agents close to agent a and ζ is a threshold. The density δ_a is calculated by counting the number of agents in a 's field of view of size w divided by the free space in that region. Notice that this region is more confined than the partially observable world ($w = 5$ in our setup). If event \mathcal{E} occurs, we consider the area as ζ -crowded; thus, the probability of deadlock significantly increases. Consequently, when an agent transitions from a non-crowded region to a ζ -crowded one, it incurs a negative reward. Conversely, when an agent moves out of a ζ -crowded region, it receives a positive reward. Moving between the non-crowded areas or ζ -crowded areas has no impact on the reward.

The threshold ζ is not fixed; its magnitude should be higher when the environment contains more agents and obstacles and lower in sparser environments. We define ζ as follows:

$$\zeta = \min\left(1.0, 0.7 + \frac{A}{m^2(1-d) - A}\right) \quad (7)$$

where A, m, d are the number of agents, grid world size, and obstacle density, respectively. The value 0.7 represents the base level of congestion tolerance. The addition of $A/(m^2(1-d) - A)$ dynamically adjusts this tolerance based on the number of agents and available space of the environment. In dense environments, the threshold becomes more lenient (i.e., higher ζ) but is capped at 1.0 to prevent it from becoming too lenient, thus maintaining a level of challenge in navigation and preventing overly dense agent clustering. We conducted tests on the lower and upper bounds within the range of 0.5 to 1.5 and determined that the range between 0.7 and 1.0 yields the best results.

Fig. 3 showcases two scenarios where the agents get a positive reward by moving out and a negative reward by moving in a ζ -crowded region. In the first scenario, consider agent number 6 (green agent) at timestep t with the field of view inside the green border; the agent is at the sparse region where the crowd density is $d = 5/13 = 0.385$. If the agent moves to the right, where the crowd density is much higher $d = 7/8 = 0.875$, the agent will be punished with a reward of -0.5 . On the contrary, in the second scenario with the agent number 8 (red agent), if the agent moves to

the left, the crowd density reduces from $d = 4/5 = 0.8$ to $d = 5/8 = 0.625$, and the agent receives a positive reward of $+0.5$. We assume the threshold $\zeta = 0.75$ in both cases.

Our reward function can be formulated as:

$$R = [R_m \quad R_c \quad R_s \quad R_e \quad R_n] \quad (8)$$

$$w = [w_m \quad w_c \quad w_s \quad w_e \quad w_n]^T \quad (9)$$

$$R_{total} = R \cdot w \quad (10)$$

In (8), $R_m = -0.3$ is the reward for the agent's movement in the four direction (N/E/S/W). $R_c = -2.0$ is the reward for agent collisions. R_s is the reward if the agent is standing still. $R_s = 0$ if the agent is at its goal, $R_s = -0.5$ otherwise. $R_e = +20.0$ is the team reward given to all agents when all agents reach their goals. R_n is the crowd-aware reward. $R_n = +0.5$ if the agent moves out of a ζ -crowded region, $R_n = -0.5$ if the agent moves in a ζ -crowded one. $R_n = 0$ otherwise. In (9), w_m, w_c, w_e, w_n are the counts of each reward type in an episode.

E. Graph neural network for local communication

We utilize Graph Neural Networks (GNNs) to facilitate local communication within a multiagent system, leveraging the intrinsic structure and relationships among nearby agents. The GNN architecture is specifically tailored to model the spatial relationships among agents by treating each agent as a node within a graph structure, $G = (V, E)$, where V denotes nodes (agents) and E represents edges (proximities between agents). The graph's adjacency matrix, A , is dynamically computed from the local view, with edge weights, A_{ij} , encoding the distances between agents. These weights control the flow of information between nodes, determining how much influence each node has on its neighbors during updates. Each node i is associated with a feature vector x^i , derived from the CNN's output, representing the agent's observed state.

The GNN architecture consists of L layers, with each layer l outputs a graph signal represented by F_l features. This input signal is fed to the next layer and then linearly transformed by the graph filters H_l to leverage the graph's structure, represented by the adjacency matrix A . The output features F_{l+1} are then passed through an activation function σ_{l+1} . Specifically, for a given layer l , the output is given by:

$$h_{l+1}^f = \sigma_{l+1} \left(\sum_{g=1}^{F_l} H_l(A) h_l^g \right)$$

Here, h_l^i is node i 's input signal ($h_0^i = x^i$), H_l are the linear operators at layer l that leverage the graph structure for processing, usually by relying solely on local neighbor exchanges and partial information about the graph.

The final layer's output, h_L^i , encapsulates aggregated neighborhood information, emphasizing the importance of spatial relationships for informed decision-making in multi-agent systems.

F. Boosted curriculum learning

We adopt a curriculum training strategy by gradually exposing the agents to increasingly complex settings over time. Initially, the agents are trained in small and sparse environments. If the agent's performance, measured by total rewards earned, does not show improvement after n consecutive episodes or after a set number of training episodes ($50k$ in our case), we escalate the complexity of the environment. Since the environment size and density are randomly sampled from a range after each training episode, we expand these ranges at each curriculum level.

$$d_l, d_h = (\min(0.05\sigma, 0.2), \min(0.1 + 0.1\sigma, 0.6))$$

$$s_l, s_h = (\min(10 + 5\sigma, 40), \min(40 + 5\sigma, 120))$$

$$n = n \times 1.5$$

Here, σ denotes the current curriculum level, and d_l, d_h, s_l , and s_h represent the ranges of obstacle density and world size to sample from.

Additionally, our approach incorporates a boosting technique. When agents exhibit poor performance in specific environments, we address this by re-sampling those environments and subsequently retraining the agents. This strategy aims to enhance the agents' proficiency in challenging scenarios by providing them with additional exposure and training opportunities in such environments.

IV. EXPERIMENTS AND RESULTS

In this section, we present our model's outcomes in comparison to the leading multi-agent reinforcement planners, evaluated across various grid world environments.

A. Experiment setup

Our experiments were conducted on randomly generated environments featuring varying obstacle densities (0, 0.1, 0.2, 0.3) and numbers of agents (8, 16, 32, 64).

Our evaluation involves testing our model on 100 distinct environments for each specific configuration, including varying numbers of agents and obstacle densities. The final results are computed as an average across these 100 environments.

B. Evaluation metrics

We evaluate the performance of our model using the following metrics:

- **Success rate:** This metric quantifies the ratio between successfully solved scenarios and the total number of tested environments. A higher success rate indicates a greater proficiency in finding solutions.
- **Makespan:** The makespan signifies the time steps required for all agents to reach their respective goals. It is equivalent to the longest path among all agents. Smaller makespan values indicate more efficient paths and, therefore, higher performance.
- **Total moves:** This metric captures the total number of non-idle actions all agents take. It provides a measure of the overall performance and coordination of the agent team.
- **Collision count:** The collision count represents the number of collisions encountered, including collisions

TABLE I

PERFORMANCE RESULTS OF THE CRAMP MODEL ASSESSED IN VARIED ENVIRONMENTS WITH DIMENSIONS 20×20 , AGENT COUNTS RANGING FROM 8 TO 64, AND OBSTACLE DENSITIES VARYING FROM 0 TO 0.3, SPECIFIED RIGHT BELOW THE METRICS. GREEN AND ORANGE COLORS INDICATE THE BEST AND SECOND BEST RESULTS. THE HYPHEN (-) SYMBOL DENOTES INVALID ENTRIES DUE TO A 0 SUCCESS RATE.

Methods	8 agents															
	Success rate				Makespan				Total moves				Collision count			
	0	0.1	0.2	0.3	0	0.1	0.2	0.3	0	0.1	0.2	0.3	0	0.1	0.2	0.3
PRIMAL	93	90	48	15	35	63	149	234	221	223	345	565	1.94	3.02	3.03	5.98
DHC	91	87	55	11	34	63	137	242	272	240	401	639	1.66	2.72	3.74	4.17
PICO	100	96	55	25	27	42	135	205	124	143	290	463	0.59	0.62	1.31	2.32
CPL	100	99	94	65	25	31	45	124	111	121	150	282	0.30	0.60	1.20	3.20
CRAMP	100	100	95	64	27	27	40	55	117	117	144	156	0.48	0.60	0.90	1.27
Methods	16 agents															
	Success rate				Makespan				Total moves				Collision count			
	0	0.1	0.2	0.3	0	0.1	0.2	0.3	0	0.1	0.2	0.3	0	0.1	0.2	0.3
PRIMAL	92	88	50	3	57	72	176	249	482	510	766	1396	6.59	8.29	11.63	17.64
DHC	94	88	48	5	54	71	169	242	477	477	822	1504	7.28	8.95	11.75	18.70
PICO	100	95	57	7	31	49	145	240	251	299	526	1292	2.98	3.98	4.96	8.00
CPL	100	95	81	22	27	41	84	213	221	249	374	780	2.80	3.70	5.30	16.10
CRAMP	100	98	83	23	30	33	45	88	236	253	275	388	2.60	3.20	3.70	5.73
Methods	32 agents															
	Success rate				Makespan				Total moves				Collision count			
	0	0.1	0.2	0.3	0	0.1	0.2	0.3	0	0.1	0.2	0.3	0	0.1	0.2	0.3
PRIMAL	92	72	9	0	54	108	245	-	958	1094	2227	-	26.23	30.48	47.27	-
DHC	92	62	3	0	49	136	243	-	957	1145	2009	-	27.04	34.79	49.29	-
PICO	100	75	19	0	38	97	225	-	551	774	1713	-	14.80	20.62	36.28	-
CPL	100	92	50	0	32	58	159	-	471	564	1032	-	11.90	17.40	30.30	-
CRAMP	100	82	34	4	37	47	85	170	489	552	754	1130	11.09	16.52	23.37	26.50
Methods	64 agents															
	Success rate				Makespan				Total moves				Collision count			
	0	0.1	0.2	0.3	0	0.1	0.2	0.3	0	0.1	0.2	0.3	0	0.1	0.2	0.3
PRIMAL	75	7	0	0	111	242	-	-	2419	3680	-	-	115.79	171.31	-	-
DHC	72	0	0	0	109	-	-	-	2121	-	-	-	106.78	-	-	-
PICO	83	13	0	0	94	225	-	-	1473	2621	-	-	90.95	128.40	-	-
CPL	80	20	0	0	92	128	-	-	1230	2204	-	-	84.00	109.00	-	-
CRAMP	78	27	2	0	90	125	278	-	1383	1877	2762	-	80.92	104.25	159.00	-

with obstacles, other agents, or instances of agents moving beyond the environment boundaries.

Regarding all of these metrics, except for the success rate, smaller values indicate better performance.

C. Result analysis

The results presented in Table I illustrate a comprehensive evaluation of our model’s capabilities in navigating complex multi-agent environments, highlighting its strengths in terms of success rates, path quality, team performance, and collision avoidance. To ensure a fair comparison with other benchmarking methods such as PRIMAL [15], DHC [39], PICO [17], and CPL [18], we maintain a fixed world size of 20×20 at test time. It is important to note that we were unable to reproduce the results reported in these papers due to implementation errors confirmed by the authors or because the code was not published. This limitation has influenced our ability to perform a direct comparison, and as such, the results for those methods are taken directly from PICO [17] and CPL [18] papers.

As depicted in Table I, our model consistently demonstrates better performance across most environments, particularly in terms of success rate. Notably, our approach can find solutions in two densely populated scenarios, featuring 32 agents with a 0.3 density and 64 agents with a 0.2 density, where all other methods encounter complete failure. Across the remaining configurations, our model consistently either secures the top position or closely follows the best-performing method.

Our model ranks among the top performers when assessing solution quality metrics such as makespan and total moves. In highly congested environments, such as those with 16 agents with a 0.3 density and 32 agents with a 0.2 density, our solutions prove to be more efficient by cutting the makespan and total moves in half.

Regarding collision avoidance, our solutions consistently yield the lowest collision counts, except for the simplest environment, where all models perform equally well. It is important to note that our collision count considers all types of collisions, in contrast to other methods that only consider collisions between agents.

To evaluate our model in larger and denser environments, we conducted experiments utilizing expanded world dimensions of 40×40 and 80×80 , maintaining a consistent obstacle density of 0.3. Comparative analysis involving CRAMP and other methods listed in Table I was performed, excluding CPL due to unavailability of their code and model. The results are presented in Fig. 4.

As illustrated in Fig. 4, CRAMP outperforms other methods in both settings. In the 40×40 environments, CRAMP attains the success rate of 93% for 4 agents, while other methods achieve approximately 85%. Although success rates decrease with an increasing number of agents across all methods, CRAMP consistently maintains its lead. A parallel trend is observed within the 80×80 environments. Despite decreased success rates due to expanded search space, CRAMP maintains a substantial lead over alternative methods.




Fig. 4. Performance comparison between CRAMP and other methods in larger and denser worlds.

D. Ablation study




Fig. 5. Effect of a crowd-aware reward function and GNN-based local communication.

To examine the impact of the crowd-aware reward function and local communication via GNNs, we conduct an ablation study. This study involves individually activating and deactivating these features to observe their influence on performance, as indicated by the episode rewards collected. The results shown in Fig. 5 demonstrate that the crowd-aware reward function significantly enhances model convergence. Specifically, models utilizing the crowd-aware reward achieve much quicker convergence, reaching approximately 120 rewards per episode, whereas models lacking this feature only converge to about -50 rewards per episode. Similarly, incorporating GNNs positively influences convergence, with models enabled with GNNs reaching 50 rewards per episode compared to those without GNNs, which fall below 0 reward.

TABLE II

EVALUATION IN 80×80 ENVIRONMENTS WITH DENSITY OF 0.1, MEASURED BY SUCCESS RATE (SR) AND MAKESPAN (MS)

Agents	W/o crowd-aware		W/o GNN		CRAMP full	
	SR	MS	SR	MS	SR	MS
4	75	99.78	79	96.54	85	67.34
8	60	143.58	65	114.32	76	98.15
16	47	280.02	50	266.47	57	225.14
32	9	543.82	12	502.12	18	407.56
64	3	913.07	5	873.59	9	746.83

We conducted additional experiments within 80×80 environments with a density of 0.1 to assess the effectiveness of our method. The findings, presented in Table II, showcase

significant enhancements in CRAMP’s performance upon the integration of crowd-aware rewards and GNN-based local communication. Enabling the crowd-aware reward function consistently boosts the success rate across various agent configurations. For instance, an increase of 16% in the success rate and a reduction of 45.43 in average makespan for 8 agents highlighting its vital role in promoting efficient navigation. Similarly, although to a slightly lesser degree, activating GNN-based local communication also yields improvements in the success rate, with a 7% increase observed for 16 agents, implying its beneficial impact on agent coordination.

V. CONCLUSION

We proposed CRAMP, a novel reinforcement learning framework for Multi-Agent Path Finding (MAPF) tasks. The innovation of CRAMP lies in the carefully designed crowd-aware reward function and our GNN-based local communication mechanism. We optimize our model using a boosted curriculum training strategy. Our extensive experiments demonstrate CRAMP’s capability by consistently outperforming other leading multi-agent reinforcement learning approaches across various metrics. Our method significantly enhances the quality of solutions in all tested configurations, making it a promising contribution to multi-agent path planning. Our approach can be applied to real-world applications, from autonomous robotics to intelligent transportation systems, where efficiency and safety are critical constraints.

Limitations: Our method may not prevent deadlocks when two robots approach the same corridor from opposite ends. Addressing such deadlocks could require alternative strategies, like prioritizing one robot’s passage over the other. Additionally, due to limited computational resources, we could only train up to eight distributed agents simultaneously, resulting in worse performance when the number of agents is large. Scaling our method for a larger number of agents may benefit from the inclusion of additional agents or dummy agents in the simulation environment.

REFERENCES

- [1] A. Maoudj and A. L. Christensen, “Decentralized multi-agent path finding in warehouse environments for fleets of mobile robots with limited communication range,” in *Swarm Intelligence: 13th International Conference, ANTS 2022, Málaga, Spain, November 2–4, 2022, Proceedings*. Berlin, Heidelberg: Springer-Verlag, 2022, p. 104–116.
- [2] L. Wen, Y. Liu, and H. Li, “Cl-mapf: Multi-agent path finding for car-like robots with kinematic and spatiotemporal constraints,” *Robotics and Autonomous Systems*, vol. 150, p. 103997, 2022.
- [3] R. Chandra and D. Manocha, “Gameplan: Game-theoretic multi-agent planning with human drivers at intersections, roundabouts, and merging,” *CoRR*, vol. abs/2109.01896, 2021.
- [4] M. Shafiq, Z. A. Ali, and E. H. Alkhamash, “A cluster-based hierarchical-approach for the path planning of swarm,” *Applied Sciences*, vol. 11, no. 15, 2021.
- [5] Z. Cheng, L. Zhao, and Z. Shi, “Decentralized multi-uav path planning based on two-layer coordinative framework for formation rendezvous,” *IEEE Access*, vol. 10, pp. 45 695–45 708, 2022.
- [6] W. Hönig, S. Kiesel, A. Tinka, J. W. Durham, and N. Ayanian, “Persistent and robust execution of mapf schedules in warehouses,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1125–1131, 2019.
- [7] D. Patel, P. Pham, and A. Bera, “Dronerf: Real-time multi-agent drone pose optimization for computing neural radiance fields,” in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 5050–5055.
- [8] K. Fransen and J. van Eekelen, “Efficient path planning for automated guided vehicles using a* (astar) algorithm incorporating turning costs in search heuristic,” *International Journal of Production Research*, vol. 61, no. 3, pp. 707–725, 2023.

- [9] C. Ferner, G. Wagner, and H. Choset, "Odrn* optimal multirobot path planning in low dimensional search spaces," in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 3854–3859.
- [10] G. Sharon, R. Stern, A. Felner, and N. Sturtevant, "Conflict-based search for optimal multi-agent path finding," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 26, no. 1, pp. 563–569, Sep. 2021.
- [11] K. Okumura, "Lacm: search-based algorithm for quick multi-agent pathfinding," in *Proceedings of the Thirty-Seventh AAAI Conference on Artificial Intelligence and Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence and Thirteenth Symposium on Educational Advances in Artificial Intelligence*, ser. AAAI'23/IAAI'23/EAAI'23. AAAI Press, 2023. [Online]. Available: <https://doi.org/10.1609/aaai.v37i10.26377>
- [12] E. Lam, P. Le Bodic, D. D. Harabor, and P. J. Stuckey, "Branch-and-cut-and-price for multi-agent pathfinding," in *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI-19*. International Joint Conferences on Artificial Intelligence Organization, 7 2019, pp. 1289–1296.
- [13] J. Li, Z. Chen, D. Harabor, P. J. Stuckey, and S. Koenig, "Anytime multi-agent path finding via large neighborhood search," in *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, Z.-H. Zhou, Ed. International Joint Conferences on Artificial Intelligence Organization, 8 2021, pp. 4127–4135, main Track.
- [14] K. Okumura, Y. Tamura, and X. Défago, "Iterative refinement for real-time multi-robot path planning," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE Press, 2021, p. 9690–9697.
- [15] G. Sartoretti, J. Kerr, Y. Shi, G. Wagner, T. Kumar, S. Koenig, and H. Choset, "Primal: Pathfinding via reinforcement and imitation multi-agent learning," *IEEE Robotics and Automation Letters*, vol. PP, pp. 1–1, 03 2019.
- [16] M. Damani, Z. Luo, E. Wenzel, and G. Sartoretti, "Primal₂: Pathfinding via reinforcement and imitation multi-agent learning - lifelong," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 2666–2673, 2021.
- [17] W. Li, H. Chen, B. Jin, W. Tan, H. Zha, and X. Wang, "Multi-agent path finding with prioritized communication learning," in *2022 International Conference on Robotics and Automation (ICRA)*, 2022, pp. 10 695–10 701.
- [18] C. Zhao, L. Zhuang, Y. Huang, and H. Liu, "Curriculum learning based multi-agent path finding for complex environments," in *2023 International Joint Conference on Neural Networks (IJCNN)*, 2023, pp. 1–8.
- [19] L. Chen, Y. Wang, Y. Mo, Z. Miao, H. Wang, M. Feng, and S. Wang, "Multiagent path finding using deep reinforcement learning coupled with hot supervision contrastive loss," *IEEE Transactions on Industrial Electronics*, vol. 70, no. 7, pp. 7032–7040, 2023.
- [20] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," in *International Conference on Learning Representations*, 2017.
- [21] D. J. Foster, D. P. Foster, N. Golowich, and A. Rakhlin, "On the complexity of multi-agent decision making: From learning in games to partial monitoring," 2023.
- [22] H. Wang and M. Rubenstein, "Walk, stop, count, and swap: Decentralized multi-agent path finding with theoretical guarantees," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1119–1126, 2020.
- [23] S. Omidshafiei, A. Agha-mohammadi, C. Amato, and J. P. How, "Decentralized control of partially observable markov decision processes using belief space macro-actions," *CoRR*, vol. abs/1502.06030, 2015.
- [24] A. Skrynnik, A. Yakovleva, V. Davydov, K. Yakovlev, and A. I. Panov, "Hybrid policy learning for multi-agent pathfinding," *IEEE Access*, vol. 9, pp. 126 034–126 047, 2021.
- [25] A. Oroojlooyjadid and D. Hajinezhad, "A review of cooperative multi-agent deep reinforcement learning," *CoRR*, vol. abs/1908.03963, 2019.
- [26] E. Boyarski, A. Felner, D. Harabor, P. J. Stuckey, L. Cohen, J. Li, and S. Koenig, "Iterative-deepening conflict-based search," in *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, C. Bessiere, Ed. International Joint Conferences on Artificial Intelligence Organization, 7 2020, pp. 4084–4090, main track.
- [27] A. Andreychuk, K. Yakovlev, E. Boyarski, and R. Stern, "Improving continuous-time conflict based search," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 13, pp. 11 220–11 227, May 2021.
- [28] H. Lee, J. Motes, M. Morales, and N. M. Amato, "Parallel hierarchical composition conflict-based search for optimal multi-agent pathfinding," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7001–7008, 2021.
- [29] J. Li, W. Ruml, and S. Koenig, "Eecbs: A bounded-suboptimal search for multi-agent path finding," in *AAAI Conference on Artificial Intelligence*, 2020.
- [30] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.
- [31] H. Zhang, M. Yao, Z. Liu, J. Li, L. Terr, S.-H. Chan, T. Kumar, and S. Koenig, "A hierarchical approach to multi-agent path finding," *Proceedings of the International Symposium on Combinatorial Search*, vol. 12, pp. 209–211, 07 2021.
- [32] G. Wagner and H. Choset, "M*: A complete multirobot path planning algorithm with performance bounds," in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2011, pp. 3260–3267.
- [33] C. Ferner, G. Wagner, and H. Choset, "Odrn* optimal multirobot path planning in low dimensional search spaces," in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 3854–3859.
- [34] T. Standley, "Finding optimal solutions to cooperative pathfinding problems," in *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, ser. AAAI'10. AAAI Press, 2010, p. 173–178.
- [35] S. LaValle, "Rapidly-exploring random trees: A new tool for path planning," *Research Report 9811*, 1998.
- [36] L. Kavvaki, P. Svestka, J.-C. Latombe, and M. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [37] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *International Journal of Robotic Research - IJRR*, vol. 30, pp. 846–894, 06 2011.
- [38] J. Li, Z. Chen, D. Harabor, P. J. Stuckey, and S. Koenig, "Mapf-Ins2: Fast repairing for multi-agent path finding via large neighborhood search," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 36, no. 9, pp. 10 256–10 265, Jun. 2022.
- [39] Z. Ma, Y. Luo, and H. Ma, "Distributed heuristic multi-agent path finding with communication," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 8699–8705.
- [40] S. Paul and J. V. Deshmukh, "Multi agent path finding using evolutionary game theory," 2022.
- [41] Q. Li, F. Gama, A. Ribeiro, and A. Prorok, "Graph neural networks for decentralized multi-robot path planning," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 11 785–11 792.
- [42] V. Suman, P. Pham, and A. Bera, "Raist: Learning risk aware traffic interactions via spatio-temporal graph convolutional networks," in *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023, pp. 5545–5550.
- [43] S. Nayak, K. Choi, W. Ding, S. Dolan, K. Gopalakrishnan, and H. Balakrishnan, "Scalable multi-agent reinforcement learning through intelligent information aggregation," in *Proceedings of the 40th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, A. Krause, E. Brunskill, K. Cho, B. Engelhardt, S. Sabato, and J. Scarlett, Eds., vol. 202. PMLR, 23–29 Jul 2023, pp. 25 817–25 833.
- [44] D. Patel, P. Pham, K. Tiwari, and A. Bera, "Dream: Decentralized reinforcement learning for exploration and efficient energy management in multi-robot systems," 2023.
- [45] J. Peng, H. Viswanath, K. Tiwari, and A. Bera, "Graph-based decentralized task allocation for multi-robot target localization," 2023.
- [46] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, "Asynchronous methods for deep reinforcement learning," in *Proceedings of The 33rd International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. F. Balcan and K. Q. Weinberger, Eds., vol. 48. New York, New York, USA: PMLR, 20–22 Jun 2016, pp. 1928–1937.
- [47] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *Proceedings of the 35th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, J. Dy and A. Krause, Eds., vol. 80. PMLR, 10–15 Jul 2018, pp. 1861–1870.
- [48] S. Schaal, "Learning from demonstration," in *Advances in Neural Information Processing Systems*, M. Mozer, M. Jordan, and T. Petsche, Eds., vol. 9. MIT Press, 1996.
- [49] A. Nair, B. McGrew, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Overcoming exploration in reinforcement learning with demonstrations," *CoRR*, vol. abs/1709.10089, 2017.
- [50] K. Pertsch, Y. Lee, Y. Wu, and J. J. Lim, "Demonstration-guided reinforcement learning with learned skills," *5th Conference on Robot Learning*, 2021.