

Decentralized Multi-Agent Path Finding Framework and Strategies Based on Automated Negotiation

Mehmet Onur Keskin

Özyeğin University

Furkan Canturk

Özyeğin University

Cihan Eran

Özyeğin University

Reyhan Aydogan (✉ reyhan.aydogan@ozyegin.edu.tr)

Delft University of Technology

Research Article

Keywords: Multi-Agent Path Finding, Negotiation, Decentralized Coordination

Posted Date: June 23rd, 2023

DOI: <https://doi.org/10.21203/rs.3.rs-3044841/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Decentralized Multi-Agent Path Finding Framework and Strategies Based on Automated Negotiation

M. Onur Keskin · Furkan Cantürk ·
Cihan Eran · Reyhan Aydoğan

the date of receipt and acceptance should be inserted later

Abstract This paper introduces a negotiation framework to solve Multi-Agent Path Finding (MAPF) Problem for self-interested agents in a decentralized fashion. The framework aims to achieve a good trade-off between the privacy of the agents and the effectiveness of solutions. Accordingly, a token-based bilateral negotiation protocol and two negotiation strategies are presented. The experimental results over four different settings of MAPF problem showed that the proposed approach could find conflict-free path solutions albeit suboptimally, especially when the search space is large and high-density, whereas Explicit Estimation Conflict-Based Search (EECBS) struggles to find optimal solutions. Besides, deploying a sophisticated negotiation strategy that utilizes information about local density for generating alternative paths can yield remarkably better solution performance in this negotiation framework.

Keywords Multi-Agent Path Finding, Negotiation, Decentralized Coordination

1 Introduction

Technological advancements in the last decades have enabled autonomous robots and vehicles to carry out various tasks, such as surveillance and transportation. They may need to navigate from one location to another to achieve their goal. Imagine an environment where hundreds of autonomous robots aim to reach specific locations. Such an environment requires a coordination mechanism to avoid some potential collisions. This problem, allocating conflict-free paths to agents to navigate safely in an environment, is well-addressed in the field of Multi-Agent Systems and known as Multi-Agent Path Finding (MAPF) problem (26; 29). Many studies tackle this problem; some propose centralized solutions while others focus on decentralized solutions (18; 28).

Computer Science, Özyeğin University, Istanbul, Turkey

E-mail: onur.keskin, furkan.canturk, cihan.eran@ozu.edu.tr, reyhan.aydogan@ozyegin.edu.tr

Centralized solutions rely on full access to all relevant information regarding the agents and properties of the given environment so that a global solution can be derived. In contrast, decentralized solutions decouple the problem into local chunks and address the conflicts locally (11). Without any time constraints, centralized approaches can find optimal solutions if there exist. However, the performance of centralized solution approaches can suffer in high-density and complex environments (9; 21; 27). Besides, complete information may not always be available due to the limitation of communication, sensors, or privacy issues. Decentralized approaches can deal with uncertainty and scalability issues and produce admissible solutions. However, they may overlook a potential optimal solution and end up with sub-optimal solutions.

Accordingly, this paper introduces a negotiation-based decentralized conflict resolution approach to MAPF in which agents autonomously negotiate with each other to refine their paths to avoid possible collisions. One of the main challenges is dealing with the uncertainty about the environment due to the limited capacity of the sensors, communication, or some privacy concerns. Most real-world applications are partially observable, where the agents can perceive some relevant aspects of the environment. For instance, a robot may not perceive all objects far from its current location. Light detection and range-finding sensors can detect up to a certain distance. Similarly, wireless communication systems also have limited communication capabilities. Agents may not exchange information with each other if they are out of a communication range. Besides, complete information is not always available due to the characteristics of the environment. For example, drivers in traffic do not know where other drivers are going. Furthermore, agents may be reluctant to share all information due to their privacy. For instance, they may not be willing to reveal their destination.

In the proposed approach, agents share their partial path information with only relevant agents close to them. If they detect any conflict on their partial path, they start a bilateral negotiation to allocate required locations for specific time steps. This study introduces a variant of alternating offers protocol enriched with token exchanges to govern this negotiation. By enforcing the usage of tokens, the protocol leads agents to act collaboratively and search for alternative optimal or sub-optimal paths to avoid conflict. Two negotiation strategies are presented in line with the protocol: Path-Aware and Heatmap. While Path-Aware negotiation strategy focuses on the agent's information on its paths and makes bids considering the length of its final path, the Heatmap strategy utilizes other agents' available path information in its field of view to avoid potential conflicts by making bids.

There are a few attempts to solve MAPF problem considering a negotiation-based approach (14; 15; 22; 23). Either they require sharing complete path information with others, consider one-step decisions such as who will move in a certain direction at the time of conflict, or aim to resolve the conflict in one shot (i.e., collaborate or reject). A detailed analysis of the proposed approach is also made in the following sections. In contrast, our approach aims to reduce the complexity of the problem by resolving the conflicts in sub-path plans iter-

actively instead of the entire path plans, thereby respecting the agents' privacy to some extent.

We evaluate the proposed approach in various settings, considering agents' ability to wait or disappear at their final destinations. We also examine the impact of three types of commitments on the agreed paths regarding the success rate and solution quality. Furthermore, we define a metric called *information sharing rates* to discuss the trade-off between privacy and solution performance for MAPF problems. Our experimental results show that our decentralized approaches can provide higher success rates in high-density scenarios (e.g., when the number of agents is 60 or 80) compared to the state-of-the-art centralized approach, namely EECBS. Although its variant, called EECBS with 0.1 suboptimality, could achieve a better solution rate and quality than the proposed approaches, our approach can achieve a better trade-off between privacy and performance. It is worth mentioning that centralized approaches entirely ignore the agents' privacy, whereas our approach prioritizes agents' privacy. Besides, we compared the performance of the proposed negotiation strategies empirically. Based on our experimental results, Heatmap agent performs better than Path-Aware agent in terms of success rate and information sharing rate across all settings. To sum up, the main contribution of this study is fourfold as follows:

- Introducing a novel negotiation strategy called Heatmap, which generates offers by avoiding dense areas to resolve the conflicts in MAPF problems,
- Presenting different commitment types and analyzing the performance of the proposed approaches in terms of various metrics,
- Studying the effect of field of view size and different problem settings where agents can take a wait action and/or disappear when they reach their destinations, and
- Investigating the trade-off between success rate and privacy empirically by examining information sharing rate and success rate of centralized and decentralized solution approaches.

The rest of the paper is organized as follows. Section 2 presents the related work on decentralized approaches for MAPF problems. While Section 3 describes the problem addressed in this paper, while Section 4 lays out the proposed negotiation approach. Our experimental setup and results are presented in section 5. Finally, this paper's main contributions and planned future work are discussed in Section 6.

2 Related Work

Conflict resolution approaches for MAPF problems can be classified according to the centralization of solution mechanisms and cooperation of agents. Centralized solution approaches path-finding of cooperative agents provide optimal plans (12). Suppose a trusted center with the information of all agents

moving in a known environment or the ability to operate all of them is unavailable. In that case, negotiation can be used for a conflict resolution mechanism (1; 13; 23; 22; 30; 14). One negotiation approach to allocating resources to multiple parties is Combinatorial Auction (CA). To resolve conflicts between self-interested agents, Amir *et al.* reduce the MAPF problem to CA and implements IBundle (20), an iterative CA algorithm (1). One crucial factor for the effectiveness of CA algorithms is the trustworthiness of self-interested agents about their utilization. Considering this issue, Amir *et al.* propose Vickrey-Clarke-Groves (VCG) auction for MAPF, a strategy-proof auction mechanism for manipulation attempts by the agents (1). In this IBundle auction, the auctioneer is exposed to a computational burden as agents submit their all-desirable bundles, which requires impractical auction time. Addressing this limitation of IBundle, the auction mechanism proposed by Gautier *et al.* can provide a feasible allocation more likely by allowing agents to submit a limited number of bundles, and so terminates in less time (13). Compared to these CA-based conflict resolutions for MAPF, our solution is free of trustworthiness and computational load concerns as it applies bilateral negotiation to resolve each conflict.

Critical challenges of addressing the MAPF problem within the decentralized method can be summarized as establishing a solution framework for agents to use while interacting with the environment, defining an interaction protocol between agents, and designing agents that can reach a solution. Purwin *et al.* propose a decentralized framework where cooperative agents allocate portions of the environment in which they can move (23). Using these frameworks, agents send information to others via a network and keep information about all. Here, agents who collide update their trajectories with a pairwise agreement based on a priority rule. Similarly, agents broadcast information in our framework while seeking a conflict-free path, resolving conflicts bilaterally. However, here revealed information is local and only achievable by agents in the broadcast range. The work of Pritchett *et al.* defines a bilateral negotiation structure to prevent collisions in air traffic (22). Each round of the negotiation session offers cost increases until an agreement is reached. This offering mechanism forces agents to concede over time. In contrast, Token-based Alternating Offers Protocol (TAOP) that we use in our study does not expose a catalyst to reduce the number of offers. It is left open to be addressed by agent strategy designs. Another bilateral negotiation approach for air traffic is proposed for the pre-flight planning of unmanned air vehicles (UAVs) (14). In this approach, agents bargain to undertake the cost of replanning to fix some of the path conflicts with counterparts in the negotiation. While this approach resolves path conflicts in pre-flight planning among all UAVs, our study addresses the real-time resolution of the MAPF problem. Besides, our negotiation approach depends on revealing sub-paths in the bilateral negotiations to resolve local conflicts so that the privacy of the agents' paths can be preserved better. Conversely, full-path information is revealed in the negotiation schema proposed by Ho *et al.*

Unlike bilateral conflict resolution approaches, Sujit *et al.* focuses on solving a task allocation problem in their work, using a multilateral negotiation structure (30). Besides, in their work, agents utilize the presented token to determine whose offer to accept in a deadlock situation that might happen, in which the agent with the least number of tokens is selected. Whereas in our framework, token usage is dedicated to penalizing repeated offers. Inotsume *et al.* demonstrates a bilateral negotiation-based approach to MAPF (15). In this approach, agents submit their desired paths to an area manager before moving their initial positions. The area manager is responsible for keeping control of a path reservation table and responses the path requests by agents based on this table to prevent any intersecting paths. The requests are accepted on a first-come, first-serve basis. Agents can request already reserved areas from corresponding agents by offering some tokens. Using a path reservation system deviates from the decentralized MAPF approach. In our framework, agents aim to minimize path costs and maximize tokens. In contrast, our solution uses tokens as an instrument to manage the negotiation behavior of any agent designed to reach an agreement.

MAPF problem requires all of the agents to reach their destination. Therefore, they must behave in a way that complies with the social welfare to some extent because allowing agents to act selfishly usually leads to a Nash Equilibrium (25). To motivate self-interested agents in a network to behave in a socially desirable way, taxation is a classical approach to diminish congestion in the network (8). Considering the MAPF problem for self-interested agents, Bnaya *et al.* propose Iterative Taxing Framework to detect where, when, and how much additional cost is imposed on agents to change their paths so that they are discouraged from using paths that would collide with others (5). Their results show a substantial total cost decrease compared to a solution scheme depending on a selfish replanning procedure. However, they do not benchmark with optimal solution cost. Thus, the efficiency of the proposed framework needs to be investigated further. Aiming the system efficiency, Ramos *et al.* introduce a toll-based multi-agent reinforcement learning approach to minimize total congestion and can achieve system-efficient results (19). In this study, tolls are used to divert drivers from congested routes. Similar to these two studies, our proposed system for MAPF aims to urge agents to prefer alternative optimal paths or individual sub-optimal paths when they face a density of other agents in their next sub-paths requiring more tokens to resolve conflicts.

Besides providing a scalable multi-agent planning framework for path-finding problems among self-interested agents, we also address the privacy aspect in the multi-agent planning domain with our study. In the privacy-concerned planning concept, agents have the task of searching and are enforced to preserve the disclosure of sensitive data while resolving conflicts of interest. MA-STRIPS (7) model provides a privacy-preserving planning scheme to control what part of the information is shared with all agents. However, this planning paradigm reveals information among all agents instead of transmitting data only between the agents subject to a conflict of interest. Bonisoli *et*

al. provides a generalized form of MA-STRIPS by establishing pairwise privacy among agents (6). In their setting, agents can specify a subset of agents intended to share information. Their approach is well proper for multi-agent problems, including individual interest protection. Similarly, in our situation, agents need to collaborate with other agents in case of future conflicts, although they are all self-interested. Therefore, path-seeker agents have to reveal local information. Our decentralized MAPF framework follows a geographical manner to condition what part of the information is shared with which agents. Such local information revealed with close agents is needed because the addressed problem requires no collision and proactive responses (i.e., replanning) by the stakeholder agents in a reasonable time before potential path conflicts.

3 Problem Definition

MAPF is the problem of assigning conflict-free paths to agents from their respective starting locations to their destinations (29). There are k agents $A = \{ A_1, A_2, \dots, A_k \}$ navigating in an undirected graph $G = (V, E)$ where starting and destination location for each agent A_i are denoted by $s_i \in V$ and $g_i \in V$, respectively. A path is a sequence of states $\pi_i = [(s_i, 0), \dots, (g_i, T)]$, where a state is a tuple of vertex $v \in V$ and time step $t \in \{0, \dots, T\}$, and T is the time step when A_i reaches to g_i . The cost of a path is calculated as $|\pi| - 1$. At any time step, agents cannot be located in the same vertex - $\pi_i^t \neq \pi_j^t \forall i \neq j$ and traverse the same edge $\pi_i^t = \pi_j^{t+1} \wedge \pi_i^{t+1} = \pi_j^t, \forall i \neq j$ (i.e., swapping conflict).

Agents are located in a $M \times M$ grid-like environment where each cell corresponds to a vertex as illustrated in Figure 1a. Initially, each agent has an optimal path to follow to reach their destination, shown by dashed lines in the grid. There are three agents, A , B , and C , whose planned paths are colored in their respective colors (i.e., red, blue, and green). All agents can only move to their vertical and horizontal adjacent neighboring cells (i.e., cardinal directions) or wait. The example includes a conflict between Agent A and Agent C at $t = 2$ in cell $(3, 3)$. They need to resolve this conflict to achieve their goals. In the classical MAPF problems, the aim of the system is to minimize the sum of individual costs $\sum_{i=1}^k |\pi_i|$ where the individual cost of each agent i corresponds to their path length denoted by $|\pi_i|$. Note that each individual aims to minimize their cost in decentralized MAPF problems.

This work examines four different MAPF problem settings. In *disappear-at-target* MAPF problem, agents can disappear from the environment when they reach their destinations. In contrast, agents stay at the destination in *stay-at-target* MAPF problem (i.e., $\pi_i^t \neq \pi_i^{t+1}$ whereas $\pi_i^t \neq g_i$). In another problem setting, agents can either wait or move. Accordingly, this work studies four different variants of MAPF problem as follows:

- Setting-1: Agents cannot wait until reaching their destinations. They stay at their target destinations and act as obstacles to other agents.

- Setting-2: Agents can wait at any time step until reaching their destinations. They stay at their target destinations and act as an obstacle for other agents.
- Setting-3: Agents cannot wait at any time step and disappear at their target destinations.
- Setting-4: Agents can wait at any time step until reaching their destination and disappear at their target destinations.

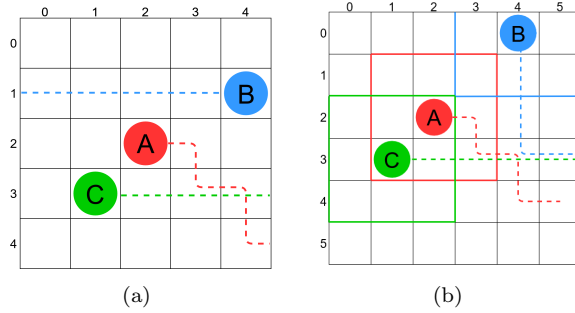


Fig. 1: Example Environment & Field of View Size = 3

4 Decentralized MAPF Framework Based on Automated Negotiation

In the proposed framework, agents are located in a grid as shown in Figure 1b. Initially, each agent knows only their starting location, destination, and a path plan to reach their destination. Those planned paths are shown as the colored dot lines in the grid for each agent. We adopt the concept of *field of view* (FoV) to simulate the partially observable environment mentioned above (10). The framework enables agents to access a limited portion of other agents' planned paths within a certain proximity and share their own. In other words, an agent's FoV determines the scope of its communication and perception capacity. An agent can only observe and communicate with other agents within its FoV. That is described as a square at size $2d + 1$, where d is the distance of agent location from the view boundary. For instance, when d equals 1, a red rectangle for Agent A shows FoV boundary. In such a case, Agent A can receive/send information from/to only agent C located in its FoV and vice versa. However, in the given snapshot, Agent B cannot communicate with other agents at that time.

In this framework, each agent broadcasts its planned subpath. Agents can determine to what extent the path will be shared with other agents in FoV. In our experiments, agents share their current subpath with a length of $2d$. If any conflict is detected by one of the agents, they can engage in a negotiation session. For example, when d equals 1, agents will share their current subpath

with a length of 2 (i.e., its next two moves) with the agents located in the scope of their FoV. Agent A broadcasts its current subpath as *Broadcast* : $[\pi_A^{t=1} = (3, 2), \pi_A^{t=2} = (3, 3)]$ while Agent C shared its own as *Broadcast* : $[\pi_C^{t=1} = (2, 3), \pi_C^{t=2} = (3, 3)]$. Since agents would detect a conflict in the vertex $(3, 3)$ at $t=2$, Agent A and Agent C start negotiating the allocation of vertices on their path.

When an agent detects a conflict with more than one agent, which negotiation to be held first is determined on *first come, first serve* fashion, and If a conflict with multiple agents is detected simultaneously, the agent to negotiate with is selected arbitrarily. For example, if d is 2, Agent B and Agent C will share their subpaths with a length of 4 with Agent A . Agent A may first negotiate with Agent B if Agent B 's message has been received before Agent C 's one. Afterward, it can encounter a bilateral negotiation with Agent C . After carrying out any successful negotiation, agents will update their path accordingly. Several negotiation sessions might be held until resolving current conflicts. If agents do not detect a conflict in their FoV, agents move to the next location in their path. Once an agent reaches its target destination, it will no longer engage in a negotiation. The negotiation is carried out according to the proposed token-based negotiation protocol. The details of this protocol and the proposed bidding strategies mainly designed for this protocol to tackle the MAPF problem are explained in the following sections.

The proposed framework requires agents to negotiate to resolve conflicts in their paths. The conflict may occur between two or multiple agents at a given time. When it happens among more than two agents, we can formulate it as multiple bilateral negotiations. As it may be harder to find a joint agreement, especially when the number of participants is high (2), the proposed approach aims to solve the conflicts in multiple consecutive bilateral negotiations. For simplicity, agents perform their bilateral negotiations consecutively. That is, a new negotiation can start after completing the previous one.

In the proposed approach, when there is a conflict in two agents' subpaths, agents negotiate to allocate the conflicted states for certain time steps. According to the example illustrated in Figure 1b, Agent A may claim to use the vertices $(3, 2)$ at $t = 1$ and $(3, 3)$ at $t = 2$ while Agent C may claim to use the vertices $(2, 3)$ at $t = 1$ and $(3, 3)$ at $t = 2$. Depending on how the negotiation proceeds, they may concede over time and change their request on vertex allocations to reach an agreement. If agents reach an agreement, they are supposed to obey the allocation for the other party. Agents are free to change their path as long as their current path allocation does not violate the agreed vertex allocation for the other party. For example, when Agent C accepts Agent A 's vertex allocation for time steps $t = 1$ and $t = 2$, Agent C confirms that it will not occupy those vertices allocated by Agent A for the agreed time steps.

The interaction between agents needs to be governed by a negotiation protocol. In automated negotiation, agents mostly follow the Stacked Alternating Offers Protocol (3), in which they exchange offers in a turn-taking fashion until reaching a predefined deadline. This protocol does not force the agents

to come up with an agreement. If both agents are selfish, they may fail the negotiation. However, having a consensus is necessary to find conflict-free paths. Therefore, agents preferably follow a protocol leading them to reach an agreement. Accordingly, we introduce a token-based negotiation protocol, namely *Token-based Alternating Offers Protocol* (TAOP) inspired by Monotonic Concession Protocol (MCP) (24) and Unmediated Single Text Protocol (USTP) (16).

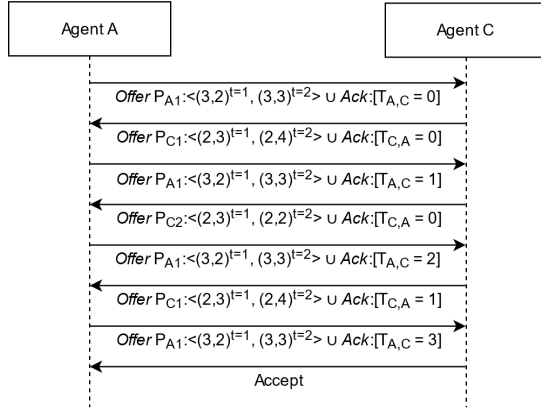


Fig. 2: Example Interaction between Negotiating Agents

According to Monotonic Concession Protocol, agents make simultaneous offers to either stick to their previous offer or make a concession. The negotiation ends without consensus (i.e., failure) if both parties stick to their previous offers. Otherwise, agents continue negotiation until reaching an agreement or fail the negotiation. This protocol leads agents to complete the negotiation without setting a predefined deadline. However, there is a risk of ending up with failure. In Unmediated Single Text Protocol, agents interchangeably become a proposer or voters during the negotiation. Initially, some tokens are given to each agent, and agents can use those tokens to override others' rejected votes. One agent starts with a random offer, and the other agent votes to accept or reject it. It is considered the most recently accepted bid if the other agent accepts. This interaction is repeated multiple times, and the most recently accepted bid is updated over time. At the end of the negotiation, the most recently accepted bid is considered the agreement. Here, the tokens incentivize truthful voting of agents to not manipulate the system by rejecting all offers. Since this protocol is particularly designed for large-scale negotiation problems, the generated bids are variants of an initial random offer, not directly applicable to our problem. On the other hand, the token idea can force the agents to concede over time in a fairway.

The proposed token-based alternating offers protocol is a variant of alternating offers protocol enriched with token exchanges. One of the agents

initiates the negotiation with an offer. The receiving party can accept this offer, make a counteroffer, or end the negotiation without agreement. The main difference is that agents cannot repeat their previous offers unless they pay for them. The protocol assumes that each agent owns a predefined number of tokens, \mathcal{T} . Those tokens enable an agent to make one of its previous offers during that negotiation. Unlike Monotonic Concession Protocol, the negotiation does not fail if both agents stick to their current offer by exchanging tokens. The essential requirement for agents is to make offers that are not previously proposed during the negotiation or pay tokens to repeat an offer. In addition to the given offer, agents send an acknowledgment message specifying the number of tokens to be used to repeat an offer previously made by the same agent. The general flow of the proposed protocol is given below:

1. One of the agents makes an offer specifying its request to allocate some vertices for specific time steps and sends an acknowledgment message regarding using its token in the current negotiation. Initially, the usage of tokens is set to zero.
2. The receiving agent can take one of the following actions:
 - ends the negotiation without any consensus.
 - accepts the received offer and complete the negotiation successfully.
 - makes an offer specifying the vertices allocation for itself that has not been offered by that agent yet and sends the acknowledgment denoting the accumulated usage of its tokens.
 - can repeat one of its previous offers, increase the usage of its tokens by one, and send the token acknowledgment message.
3. If the agent accepts or ends the negotiation, the negotiation is finished. The accepting agent receives tokens amounting to the calculated token usage difference from its opponent, $\min(\mathcal{T}_{opp} - \mathcal{T}_{self}, 0)$ where \mathcal{T}_{opp} and \mathcal{T}_{self} denote the total number of tokens used by the opponent and the accepting agent during the entire negotiation respectively. If the accepting agents spend more tokens than their opponent, it does not receive any tokens. Otherwise, the receiving agent can take any action mentioned in Step 2.

Considering the scenario given in Figure 1b, an example negotiation trace between Agent A and C is illustrated in Figure 2. Agent A initiates the negotiation with its offer P_{A1} requesting to claim the vertices $(3, 2)$ for $t = 1$ and $(3, 3)$ for $t = 2$. Agent C does not accept this offer and makes its offer specifying its allocation, such as $(2, 3)$ and $(2, 4)$. Since Agent A insists on its previous offer, it increases its token usage by one. As the example shows, agents send an acknowledgment message and offer each turn. Agent C accepts Agent A 's offer in the fourth round. It confirms that Agent C will not move to $(3, 2)^{t=1}$ and $(3, 3)^{t=2}$. In return, Agent A will pay 2 tokens ($\mathcal{T}_A - \mathcal{T}_C$). It is worth noting that the token exchange is performed at the end of the negotiation, depending on who accepts the offer. The agreement is not committed if an agent needs to pay tokens but has insufficient tokens (i.e., negotiation fails).

4.1 Commitments of Agreements

After reaching an agreement by following the proposed negotiation protocol above, agents may adhere to or retract from the agreed path allocation depending on their role. If the agent pays a token and convinces another party to accept its offer, it is free to change its remaining path later on. Otherwise, the agent accepting the offer should determine another subpath that does not conflict with the agreed allocation (i.e., the tuple of opponent agent ID, timestep, and location) and broadcast its new subpath within FoV. However, that agent may end up with suboptimal or unsolvable positions at any time step after the agreement time t' . For such situations, we define different commitment types for agents as follows:

- **Standard Commitment (SC):** The agent agreed on the offer and cannot enter the subpath allocated to the other agent according to the agreement (10).
- **Zero Commitment (ZC):** At any time step after t' , the agent can decommit the current agreement when it observes new conflicts.
- **Dynamic Commitment (DC):** At the end of the negotiation, the agent who accepted the opponent's offer should obey the agreed subpath until the conflicted state. Afterward, it can break the commitment if deemed necessary as a result of new observations at any moment after t' .

To illustrate how these commitment types work, consider the example given in Figure 1b, where the FoV size is three. Assume that Agent C gives the path advantage – $((3,3)$ position is allocated for Agent A in the grid to Agent A at the end of the negotiation at time step t .

- While ZC is utilized, Agent C can again move ahead of the Agent A path in $t + 1$ if necessary.
- While SC is utilized, when Agent C detects that there are two different conflicts with both Agent A and B and agrees to give both full path superiority, C cannot enter the paths where they are placed behind the conflicted positions. Since C already gives all usage rights of the paths in t , even if C needs to break its commitment to reach its destination.
- While DC is utilized, Agent C cannot go from Agent A 's path until it passes a conflicted point (e.g., $(3,3)$); it has to stick two steps to its commitment until $t+2$. However, suppose there is future negotiation with Agent B at $t+s$ for a conflicted location at $t+s+1$. In that case, Agent C should stick only one step to its commitment to Agent B . Agent C has a different commitment deadline for each negotiation according to its conflicted places in dynamic commitment.

For these effects, the commitment type can be a distinguishing feature for decentralized MAPF problems, especially in scenarios where the FoV is large in dense environments. Hence, this paper analyzed the effects of the commitment type as an independent variable. Due to the nature of different commitment types, agents can increase uncertain behaviors in the environment according

to the changes in their observations. These behaviors cause to increase in the number of negotiations for each scenario. In other respects, these behaviors can increase agents' individual utility and the successful completion rate of each scenario since they will increase agents' resiliency and adaptation skills. Section 5.4 will further examine the impact of these commitment differences across all other configuration variables.

4.2 Proposed Negotiation Strategies

We introduce two bidding strategies, namely *Path-Aware* (10) and *Heatmap*. Note that the agents are called by the name of the bidding strategy they adopt. While Path-Aware agent considers its available path and token information, Heatmap agent aims to reduce potential conflicts and prefers low-density routes, as explained below.

4.2.1 Path-Aware Negotiation Strategy

Existing negotiation strategies focus on only which offer to make at a given time and when to accept a given offer (4). Therefore, there is a need to design a new strategy considering token exchanges. Hereby, we propose a negotiation strategy, which we call Path-aware agent (10), determining when to repeat an offer or generate a new offer. It aims to utilize the information available to determine when to insist on its current path. It is worth noting that each agent generates possible paths leading them to its destination by using the A^* Algorithm by considering the non-conflicting paths with those of other agents in its FoV. Those possible paths constitute the agent's bid space. Afterward, they sort those paths in descending order of their cost.

Algorithm 1 describes how an agent negotiates according to Path-Aware agent. At the beginning of the negotiation, the current path in FoV (P_c) is the relevant part of the optimal path (i.e., the shortest path to its destination). It corresponds to the first offer in the negotiation. When the agent receives an offer from its opponents, it checks whether it is possible to generate a path of equal length or shorter than its current path to the destination (Line 1). If so, it accepts its opponent's offer (Line 2). Note that the path generation function takes the opponent's offer, O_{opp} , as a constraint while generating its best possible path to the destination. If the best possible route, which is in line with the opponent's bid, is better than or equal to the current path, then the agent accepts the opponent's offer. If the agent's remaining tokens are greater than the length of the remaining path to the destination (Line 4), it decides to repeat its previous offer. It updates its remaining tokens accordingly (Line 5). Recall that for each repetition, the agent needs to use one token. Otherwise, it concedes and sets the next possible best path from the sorted path space P_{Space} as its current path in its FoV (Line 7). Accordingly, the agent offers its current path in FoV P_c (Line 9). Note that agents must concede if they do not have any tokens left.

Algorithm 1: Negotiation Strategy of Path-Aware Agent

```

Data:
 $T_r, P_r$  : Agent's remaining tokens & path to destination
 $P_c$  : Current path in FoV
 $P_{Space}$  : Sorted path space
 $O_{opp}$  : Opponent offer
1 if  $|P_c| \geq |generatePath(O_{opp})|$  then
2   |  $accept()$ 
3 else
4   | if  $T_r > |P_r|$  then
5     |    $T_r --$ 
6   | else
7     |    $P_c \leftarrow P_{Space}.next()$ 
8   | end
9   |  $offer(P_c)$ 
10 end

```

4.2.2 Heatmap Negotiation Strategy

The Path-Aware agent focuses on the agent's remaining path length and the number of tokens. The density in the observable environment is neglected during the bidding process. Thus, the agreed paths may lead to further conflicts. Heatmap bidding strategy incorporates this environmental information into the evaluation of bids. This strategy uses the available data in FoV to generate a heat map of the terrain around the agent and detects high-density areas. Consequently, it avoids generating a path through dense regions as the cars avoid potential traffic jams by preferring the low-density routes in the traffic. In other words, by analyzing the received path information from other agents, the agent aims to foresee moving to which region of the environment causes path conflicts to be less likely. The generated heatmap is built based on the broadcasted information of other agents within Fov. It considers each heatmap while calculating the estimated cost of any path. Figure 3 shows an example of a heatmap created for an agent in the middle of the grid. When the location information of an agent is detected by other agents, they compute a heatmap centered at the location of the detected agent, as seen in Figure 3a. Here, the values are normalized concerning FoV (see Figure 3b).

To better understand how the agent uses a heatmap in its bidding strategy, we illustrate a toy example depicted in Figure 4. In this example, Agent *A* and Agent *B* start negotiation because they detect a swapping conflict on their paths. Therefore, when Agent *A* estimates the cost of each possible path, it considers only the heatmap of Agent *C*. As seen from Figure 4, Agent *A* generates the heatmap of Agent *C* for $t=1$ and $t=2$ based on the broadcasted path information by Agent *C*. Let us assume it will calculate the estimated cost of its initial path, $(2, 2)$, $(2, 3)$, $(3, 3)$. At $t=1$, it is supposed to be at $(2, 3)$, and the value of this cell is 0.33, while at $t=2$ it will be $(3, 3)$, and its heatmap value is also 0.33 according to the heatmap of the Agent *C*. The path cost would be estimated as $(0.66 = 0.33 + 0.33)$.

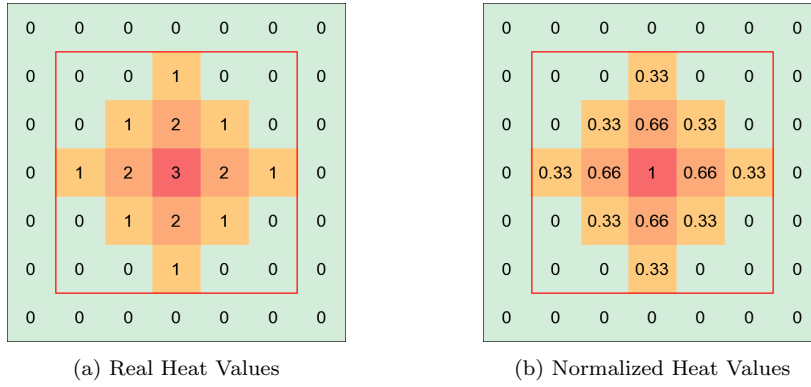


Fig. 3: Heat Map Generated by an Agent




Fig. 4: Heatmap of Agent C According to Agent A

When multiple agents are present in the heatmap, the values of the cells are added up, which is illustrated in Figure 5 depicting a portion of the heatmap generated by Agent *A* at time step t for its negotiation session with Agent *D*. It is assumed that Agent *A* can listen to broadcasts of Agent *B*, *C*, and *D*, and Agent *A* is engaged in a negotiation session with Agent *D*. Note that the heat values of the opponent agent are not considered during the negotiation. Consequently, heat values generated by Agent *D* are ignored. However, the paths of Agent *B* and *C* are assumed to be steady for the duration of the negotiation. Therefore, their broadcasted information is used to generate a heatmap, which has a portion of it seen in the figure.

Heatmap agent follows almost the same strategy as the Path-Aware strategy but takes the potential future conflicts into account; therefore, its cost estimation differs from Path-Aware agent. Unlike Path-Aware agent, which uses the actual cost of the paths, Heatmap agent estimates the cost of the bids (i.e., paths) according to Algorithm 2. The overall cost is the sum of the estimated cost of each cell in the given path. That is, the cumulative cost is estimated by considering the relevant cell value on the Heatmap of each neigh-




Fig. 5: Heat Generated by Agents B and C for Agent A

bor agent other than the current opponent agent. Line 2 iterates through the points in the given path whose cost is estimated, while Line 3 iterates through the heatmaps of the neighbor agents except for the opponent. The total cost is increased by each relevant cell value of the heatmaps (Line 4). After calculating the estimated cost of each possible bid, Heatmap agent sorts them in descending order and employs the approach explained in Algorithm 1.

Algorithm 2: Path Cost Estimation from Heatmap

Data: H_T^j : Heatmap for Agent j at a given time T N_c : Neighbor agents excluding the opponent in FoV P_c : A given path whose cost will be estimated**Output:** EC : Estimated cost of a given path, P_c

```

1  $EC \leftarrow 0$ 
2 for  $i \leftarrow 1$  to  $n \in |P_c|$  do
3   foreach  $j \in N_c$  do
4      $EC \leftarrow EC + H_{t+i}^j.get(P_c^i)$ 
5   end
6 end
7 return  $EC$ 

```

Throughout its life cycle in the environment, an agent can continuously observe its environment. By observing the changes in their environments, agents can deduce congestion in an area and behave to avoid dense regions. One way to introduce congestion avoidance while designing the agent is to introduce a system to help the agent avoid bids that would direct it toward congested regions. In this strategy, the agent utilizes the information available in FoV to generate a simple heatmap to perceive the crowd and avoid making bids leading to those regions. The generated heatmap will allow the agent to differ-

entiate between offers by weighting bids that lead into crowded regions heavier than those that do not.

5 Evaluation

To evaluate the proposed negotiation approaches, we simulate a set of scenarios for the four introduced MAPF settings. We report the empirical results in terms of several metrics: solution rate, solution quality, number of negotiations, and information sharing rate. Considering these metrics, we analyze the effect of negotiation strategy, the field of view (FoV) size, and the commitment types on the system behavior. Besides, we compare the performance of our decentralized MAPF framework with the centralized MAPF approach. In the following sections, we first introduce experimental setting and evaluation metrics. Then, we present and analyze the results.

5.1 Experimental Setting

We deploy the state-of-the-art optimal search algorithm for centralized MAPF, Conflict-based Search (CBS) citeSharon2012, to examine the performance of our decentralized MAPF approach. CBS is a two-level search algorithm in which high-level search resolves path conflicts by generating a set of constraints for conflicted agents while a single path under these constraints is generated by an optimal shortest-path algorithm, like A*, in the low-level search. Here, a constraint is a tuple (i, v, t) where agent i is prohibited from occupying vertex v at time step t . At the high-level search, a binary tree of constraint nodes is operated to resolve conflicts between paths where each node consists of paths and a set of constraints for each agent. When a conflict between two paths is detected, two child nodes are generated by constraining respective agents. In each child node, one agent in the conflict is prohibited from using a conflicted vertex or edge by adding a constraint, and a new path is searched for that agent at the low level under the new constraint set. CBS was used for benchmarking in our previous work (10), demonstrating the limited scalability of CBS.

In this work, our experiments utilize the most enhanced version of CBS, Explicit Estimation Conflict-Based Search (EECBS) (17). EECBS is a bounded sub-optimal MAPF algorithm that guarantees the solution cost by an optimality factor of w . It applies different node selection rules than CBS at low and high-level searches. Instead of a best-first search in A* at the low level, it follows a focal search in which solutions in the focal list have a cost value at most $w \cdot f$ where f is the lower bound to the optimal solution cost found during the search. EECBS manages its high-level search to reach a suboptimal solution faster by utilizing inadmissible information about potential cost increments (31). We run optimal EECBS and EECBS-1.1, which has the upper bound of 1.1 for suboptimality. We set the time limit to five minutes for

EECBS and one minute for EECBS-1.1 per scenario. For optimality evaluation, we only consider the scenarios in which the optimal solution is found with EECBS. Each setting of MAPF problem (i.e., settings 1–4) experimented with 100 different scenarios in a 16×16 grid without obstacles for 20, 40, 60, and 80-agent scenarios in standard commitment protocol. All scenarios were generated with randomly distributed initial path lengths between 4 and 24 using MAPF benchmark dataset provided by (28). Each combination of the aforementioned configurations was repeated for FoV sizes 5, 7, and 9 (i.e., FoV5, FoV7, and FoV9). For the robustness of the evaluation, these runs were repeated five times, considering randomness in the agent’s suboptimal bid selection and bilateral partner selection. All experiments were conducted on 12 computers with an 8-Core 3.2 GHz processor and 32 GB RAM.

We evaluate the performance of the proposed decentralized approaches in each problem setting by comparing their performance with that of EECBS (i.e., centralized approach) in terms of several metrics. The first metric we use is the *solution rate*, which denotes the ratio of the number of conflict-free path solutions for all agents over a given MAPF configuration. In addition to solution rate, solution quality is another criterion for benchmarking suboptimal solvers, typically defined as an *optimality gap*. For MAPF problem, optimality gap describes the rate of the path cost difference between the solutions obtained by the given approach and an optimal solver. It is formulated in Equation 1 where $|\pi_{solution}|$ is the average of the final path lengths of agents provided by the proposed approach and $|\pi_{solution}^*|$ is the average of optimal path lengths provided by EECBS. Since proposed decentralized approaches could yield solutions for some scenarios that EECBS cannot solve, we use *normalized path difference* metric to compare the performance of those decentralized approaches. It is defined in Equation 2 where N denotes the number of scenarios, $|\pi_{Dsolution}|$ is the average of the final path lengths of agents provided by the proposed approach on the i^{th} scenario and $|\pi_{Dsolution}^*|$ is the average length of the shortest paths provided by available decentralized approaches.

$$\text{Optimality Gap} = \frac{|\pi_{solution}| - |\pi_{solution}^*|}{|\pi_{solution}^*|} \quad (1)$$

$$\text{Normalized Path Difference} = \frac{1}{N} * \sum_{i=1}^N \frac{|\pi_{Dsolution}^i| - |\pi_{Dsolution}^{*i}|}{|\pi_{Dsolution}^{*i}|} \quad (2)$$

Moreover, we evaluate the proposed decentralized approaches considering the average number of negotiations and privacy of agents’ paths. We define information sharing rate (IS) to consider what extent of an agent’s final path is shared with other $k-1$ agents. Equation 3 defines the information sharing rate of agent j by averaging the rate of agent j ’s shared path with other agents where $\pi_{solution}^j$ and $\pi_{broadcasted}^{j \rightarrow i}$ denote the final path of agent j and the part of that path shared with other agents during broadcasting, respectively. For evaluating a given scenario, the information sharing rate is averaged over all agents in that scenario.

$$IS_j = \frac{1}{k-1} * \sum_{i=1}^{k-1} \frac{|\pi_{broadcasted}^{j \rightarrow i} \cap \pi_{solution}^j|}{|\pi_{solution}^j|}, i \neq j \quad (3)$$

For all metrics mentioned here, we only consider the results of the runs where each strategy could find a solution for the given scenario to evaluate the strategies fairly. The results used for each metric calculation are averaged over methods that all compared approaches could solve.

In the following sections, we first compare proposed decentralized approaches with centralized ones (Section 5.2) and analyze the decentralized MAPF approaches elaborately (Section 5.3). Afterward, we investigate the effect of the commitment types (Section 5.4).

5.2 Decentralized versus Centralized MAPF Approaches

To compare the performance of our decentralized approaches with the centralized MAPF approach, we have run EECBS and its variant EECBS-1.1 to find optimal and suboptimal solutions with an optimality gap of 0.1, respectively, for all scenarios per each set consisting of 100 scenarios. We run our proposed decentralized approaches with the standard commitment type, namely Heatmap and Path-Aware. Note that FoV size for those approaches is set to five¹.

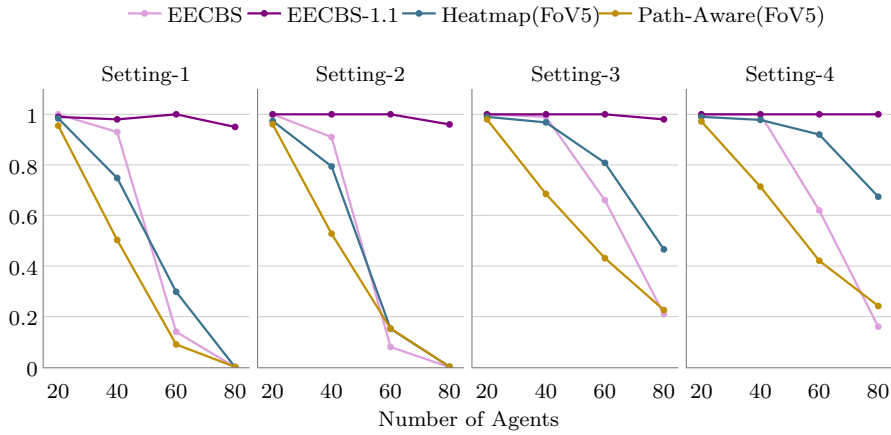


Fig. 6: Success Rates by EECBS, EECBS-1.1, Path-Aware, and Heatmap Agents with FoV(5)

Figure 6 shows the solution rate for EECBS, EECBS-1.1, Path-Aware, and Heatmap agents with the FoV size of five. It is observed that the decentralized

¹ We have tested with varying FoV values and got a higher solution rate on average when their FoV size is five (FoV5).

agents can provide higher success rates in high-density scenarios (e.g., when the number of agents is 60 or 80) compared to EECBS. Recall that EECBS may not find optimal solutions for some scenarios. On the other hand, when EECBS is allowed to perform with the suboptimality factor of 0.1, it can provide solutions for almost all scenarios.

When we investigate the scenarios for which all approaches found a solution, the success rates of EECBS, Heatmap, and Path-Aware approaches are given in each cell in Table 1 where k denotes the number of agents. It is seen that when the number of agents is equal to 20, EECBS could find optimal solutions for all scenarios. However, when the number of agents increases, it has difficulty producing a solution (see success rates when $k=80$). In comparison to its success rates, decentralized approaches, particularly Heatmap, reach higher success rates (e.g., 0.16 versus 0.67 for Setting-4)².

Table 1: The Success Rates by EECBS, Heatmap, and Path-Aware based on individually solved scenarios, respectively.

	k=20	k=40	k=60	k=80
Setting-1	1.0/0.98/0.96	0.93 /0.75/0.5	0.14/ 0.30 /0.09	0.00/0.00/0.00
Setting-2	1.0/0.97/0.96	0.91 /0.79/0.53	0.08/ 0.15 /0.15	0.00/0.00/0.00
Setting-3	1.0/0.99/0.98	0.99 /0.97/0.68	0.66/ 0.81 /0.43	0.21/ 0.47 /0.23
Setting-4	1.0/0.99/0.97	1.0/0.98/0.71	0.62/ 0.92 /0.42	0.16/ 0.67 /0.24

The decentralized approaches perform with 16% optimality gap on average, whereas EECBC-1.1 produces solutions with 6% average optimality gap compared to EECBS (Figure 7). Moreover, we can see that extending Path-Aware agent and developing Heatmap agent has improved the overall solution rate and slightly decreased the optimality gap in decentralized setup. Since Heatmap agent is designed to be aware of the environment, these results show that a more sophisticated agent design can yield better performance in the proposed negotiation-based MAPF framework.

Lastly, we investigate to what extent the decentralized approaches take into account privacy. It is worth noting that the centralized approaches may find better solutions; however, they use global information on all agents' paths to resolve conflicts, whereas agents utilize local information to find conflict-free paths in the proposed decentralized MAPF framework. Figure 8 shows the average information share rate for Path-Aware and Heatmap agents for each setting with varying FoV sizes. It is seen that when agents employ the Heatmap strategy during their negotiation, the information shared with other agents is significantly lower than that of Path-Aware agents. Since Heatmap agent employs a density-aware strategy, which avoids crowded areas, it has a tendency to change its path rapidly.

Moreover, it is observed that the average information sharing rate is higher when FoV is larger irrespective of the condition, as is expected since a larger

² Those results are boldfaced.

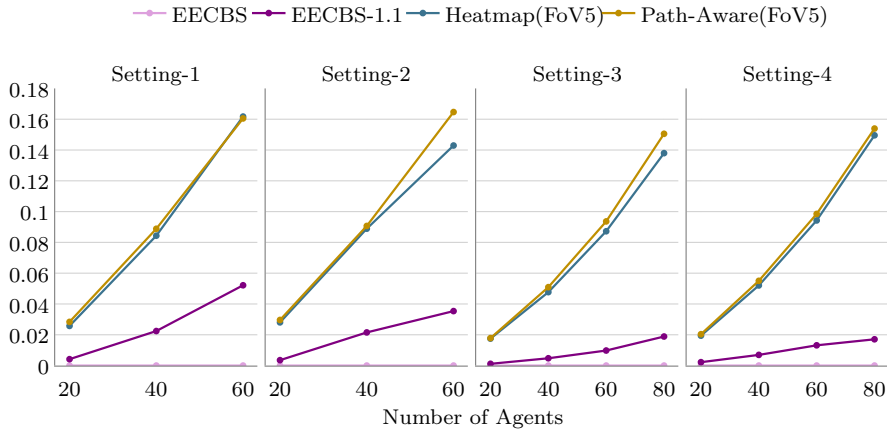


Fig. 7: Optimality Gaps by EECBS, EECBS-1.1, Path-Aware, and Heatmap Agents with FoV(5)

FoV makes agents perceive more neighbor agents. Contrary to expectations, when the environment becomes denser (i.e., involving a higher number of agents), the information sharing rate and the number of broadcasted agents decrease. Our observation is that the percentage of agents encountered in FoV is increased relatively less than the growth rate of the population (see Figure 9). Therefore, even though the number of message exchanges is similar in both settings, as seen in Figure 9, the proportion of its final path is shared with others relatively less than the case of the Path-aware agents.




Fig. 8: Average Information Sharing Rate for Each Agent




Fig. 9: The Ratio of The Average Number of Broadcasted Agents to Total Number of Agents

Besides, we might expect to see a decrease in the information sharing rate when dealing with a MAPF problem in the case of the agents disappearing at their targets (Settings 3 & 4) because they are not obstacles for others anymore when they reach their destination. However, when we analyzed the scenarios solved, we noticed that agents mostly solved less conflicting scenarios in all settings, and agents in settings 3–4 could also solve more complicated scenarios. Therefore, agents in those settings encountered more agents; thus, the average information sharing rate is higher in those settings.

5.3 Analysis of Decentralized MAPF Approaches

As discussed before, Heatmap agents are capable of providing solutions in challenging settings where agents have to take action in each time step and act as an obstacle for other agents in the system when they reach their destinations (i.e., standard setting – setting 1). This section examines how the proposed negotiation approaches handle the same scenarios if we enable agents to wait at any time (i.e., settings 2–4) and disappear from the environment when they reach their goals (i.e., settings 3–4). Consequently, we evaluate the performance of the proposed approaches on the previously generated 16×16 scenarios for four different settings of MAPF problem as defined in Section 3.

Figure 10 presents the success rates of Heatmap and Path-Aware agents in each environment density with each FoV size (i.e., $\text{FoV} = 5, 7, \text{ and } 9$) for all problem settings. Enabling wait action and the behavior of disappearing at the destination (Setting 4) has a positive effect on the solution rates. Both agent types perform much better when disappearing at target, especially for scenarios consisting of 40, 60, and 80 agents. That may stem from the fact that the complexity of the problem is reduced since the agents at targets are

no longer obstacles for others. Another observation is that the performance of Heatmap agent is affected more positively when the agents are allowed to wait anytime, especially in the case of the disappearing setting (Setting-3 versus Setting-4).




Fig. 10: Success Rates by Heatmap and Path-Aware Agents for Each FoV Size

The solution rate drastically drops when we increase the number of agents to 80, especially when the agents are allowed to take wait-action to resolve path conflicts with others. The empty cells in the 16x16 grids significantly decrease when there are 80 agents. In addition, allowing agents to wait when they can act as an obstacle for other agents. To sum up, this will lead to more conflicts among agents, which cannot be straightforwardly resolved. Most of the cases, agents are stuck and cannot move in any direction – ending up with a low solution rate.

Regarding the quality of the solutions, Figure 11 shows the average normalized path differences for both agent types in four settings with varying FoV sizes. For 80-agent scenarios of Settings 1 and 2, average normalized path differences are missing since the solution rate is equal to zero in these problem configurations. It is observed that solution quality differences between Path-aware and Heatmap agents, Heatmap agent is slightly better than Path-aware while the success rates of Heatmap agent are higher. That shows the superiority of Heatmap agent over Path-aware agent. Heatmap agents can resolve path-conflicts without sacrificing path length. Heatmap agent can achieve solutions at slightly lower average normalized path differences than Path-Aware agent overall. This result promotes the design of environment-aware agents for MAPF problems rather than agent designs only focusing on token consumption. Figure 12 shows the average number of negotiations per agent. In general, agents have an almost equal number of negotiations in each setting except when the number of agents is 80. In that case, the higher FoV is, the higher number of negotiations occur (see setting 4).




Fig. 11: Average Normalized Path Differences for Heatmap and Path-Aware Agents for Each FoV Size




Fig. 12: Average Number of Negotiations for Each Agent Type

5.4 The Effect of Commitment Types

In this section, we provide a focal analysis of the system performance using the proposed negotiation approach to examine how it evolves in dense environments; therefore, we consider only 80-agent scenarios for Setting-3 and Setting-4 MAPF configurations. Recall that agents disappear when they reach their target in both settings. Agents can take a wait action at any time step until reaching their destination in Setting 4, whereas they cannot wait in Setting 3. Generally, Heatmap agent has provided better results than Path-aware agent in many metrics when using the standard commitment protocol. Hence,

we only focus on Heatmap agent to analyze the effect of the commitment types. All results throughout the section are provided in the taxonomy of agent configurations, which are the commitment types (SC, ZC, and DC) and the FoV sizes (5, 7, and 9). We consider the aforementioned performance metrics other than the optimality gap since EECBS could not provide optimal solutions for the scenarios involving 80 agents.

Depending on their commitment types, agents may adhere to or retract from the agreed path allocation from their previous negotiations. While agents utilizing SC entirely stick to their prior agreements and generate offers that do not conflict with the previous agreements, agents utilizing ZC can change their path and conflict with their previous commitments at any time after the agreement time. In contrast, agents utilizing DC must stick to the agreed allocation until the time step of resolved conflict (i.e., previously conflicted state), but after that state, the agent is free to change its path. Figure 13 shows the success rates of Heatmap agent for setting 3 and 4. When we analyze those results, it is seen that Heatmap agent reaches more solutions in zero commitment than SC in both settings. In particular, the success rate differences between zero and standard commitments are %31-%35-%52 in Setting-3 and %18-%18-%40 in Setting-4 with the respective FoV sizes. In SC, the agents narrow their search space by sticking to their previous agreements; thus, this situation may prevent agents from finding an agreement in further negotiations. As a result, allowing agents to decommit their previous negotiation agreements may increase the chance of reaching their final destination. In addition, the success rate differences between agents utilizing ZC and DC are minimal (%1-%1-%0 in Setting-3 and %6-%1-%3 in Setting-4 with the respective FoV sizes). It is worth noting that the success rates of agents utilizing ZC and DC slightly increase, whereas the success rates of agents utilizing SC drastically decrease when FOV size increases.

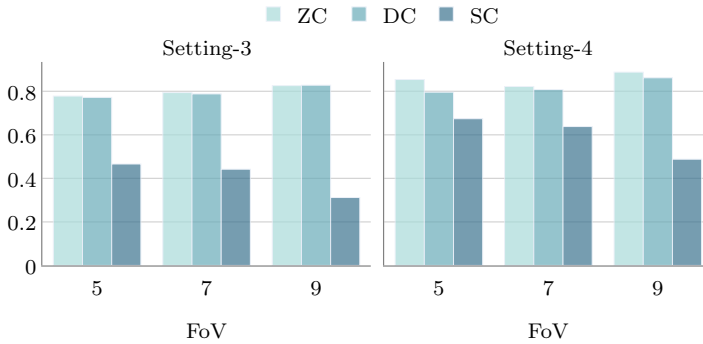


Fig. 13: Success Rates for Each Commitment Type

As seen from Figure 14, there is no evidence of remarkable differentiation between medians (also deviations) of the normalized path difference in Setting-

3 and Setting-4. When we apply the Mann-Whitney U statistical tests, there is no statistically significant difference among SC, ZC, and DC ($p > 0.05$) except in the case when the FOV is equal to 5 in Setting 3 (SC vs. ZC: $p = 0.046 < 0.05$) and SC vs. DC: $p = 0.024 < 0.05$). This supports the observation that agents utilizing ZC and DC can accomplish finding more non-conflicting paths for the given scenarios while their solutions are as cost-efficient as in the case of the SC.

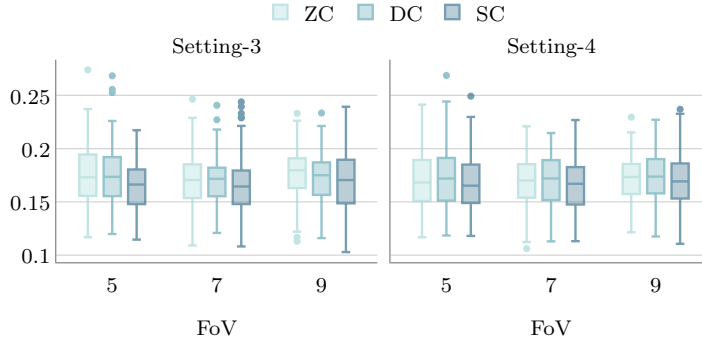


Fig. 14: Normalized Path Difference for Each Commitment Type

Moreover, we present a new metric to analyze how often agents interact with each other to resolve observed conflicts throughout a scenario in our negotiation-based decentralized MAPF system for each commitment type. The average number of negotiations per time step is presented for each commitment type and each FoV size in Figure 15. This metric provides granular information about the interaction level of agents in a decentralized system. Agents engaged in a comparably higher number of negotiations in the system's initial state (i.e., $t=0$), and the number of negotiations significantly decreased at $t=1$ in all FoV configurations. Note that the average number of negotiations naturally decreases as the number of agents who arrive own destination increases.




Fig. 15: Average Number of Negotiations per Time Step in Setting-4

Secondly, another distinctive effect of the commitment types comes from the system behavior differences concerning FoV size. Agents have a chance to resolve more number conflicts beforehand within a time step with larger FoV sizes. Hence, agents engage less frequently with FoV9, continuously decreasing the number of negotiations. On the other hand, the average number of negotiations between $t= 1-4$ increases in the FoV5 setting and remains stable in the FoV7 setting. This characteristic curve change from FoV size 5 to 9 reveals how information sharing among self-interested pathfinder agents affects the conflict resolution behavior of a decentralized system. Agents with narrow FoV observe less number of conflicts beforehand but have to repeatedly interact with near agents to resolve conflicts at earlier stages. On the other hand, agents with larger FoV can resolve more conflicts at once and monotonically fewer interactions with each other later. It indicates that the system converges to a solution faster with larger FoVs, and the solution rate slightly rises without compensation of solution quality as seen in Figure 13 and 14.

We also consider the privacy of self-interested agents as a third dimension for evaluating a decentralized MAPF system performance alongside success rate and solution quality metrics. Information sharing among the agents is significantly higher when they employ ZC and DC commitment types than SC, as seen in Figure 16. Another observation is that the increase in the information sharing rate is higher for ZC and DC when the FoV size of agents increases since the agents with SC engage in less number of negotiations (see Figure 15), so they share less unique information about their planned paths with others.

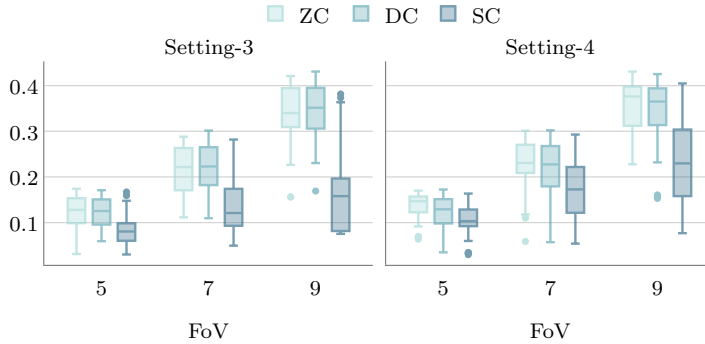


Fig. 16: Average Information Sharing Rates for Each Commitment Type

On the other hand, Heatmap agents with ZC and DC share more information to resolve conflicts, thus achieving much higher success rates (see Figure 13). The success rate and the solution quality are evaluated proportionally with the information sharing rate, which is presented in Figure 17 and 18, respectively. The system performance trades in favor of privacy by a clear margin when agents have a narrow FoV. Note that trade-offs between privacy and solution performance in terms of success rate and quality are linearly

taken. Compared to the FoV5 configuration, the success rate over the average information sharing rate substantially decreases in Setting 3 and 4 with FoV7 and FoV9 configurations. Similarly, the decrease in the solution quality over the average information sharing rate is significant when agents have FoV7 and FoV9 in both problem settings.




Fig. 17: The Ratio of Success Rate to Average Information Sharing Rate for Each Commitment Type (Higher is better.)

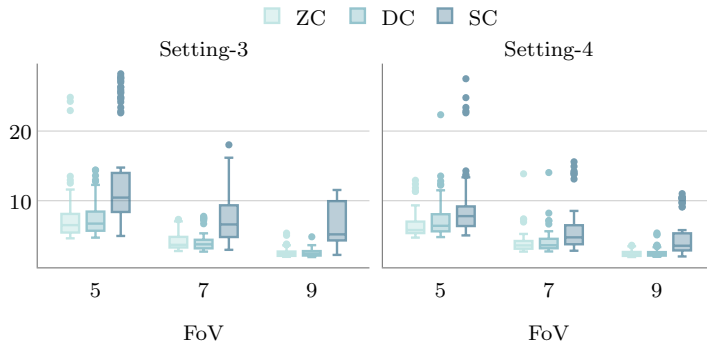


Fig. 18: The Ratio of '1 - Normalized Path Difference' to Average Information Sharing Rate for Each Commitment Type (Higher is better)

6 Conclusion

This paper addresses how self-interested agents can coordinate in a grid environment to reach their destination without collision and propose a negotiation approach to resolve conflicts between their paths. Accordingly, we propose a negotiation protocol and two compatible negotiation strategies: Path-Aware and Heatmap. While the Path-Aware agent is more concerned about the length of its absolute path, Heatmap agent considers future conflict among agents.

Heatmap agent exploits the broadcasted path information by other agents accordingly. The proposed approach is evaluated elaborately in different settings where agents are allowed/disallowed to take a wait action and/or disappear when they reach their final destinations. In addition, we introduce three types of commitments and study their effect on the success rate and solution quality. Furthermore, we present a new metric called *information sharing rates* to discuss the trade-off between privacy and solution performances for MAPF problems.

According to the experimental result, Heatmap agent outperforms Path-Aware agent concerning the success rate and information sharing rate in all settings. Moreover, Heatmap agent achieves higher success rates than the optimal algorithm for MAPF, EECBS, when agents can take wait action and disappear at their targets (Setting 3 and 4). On the other hand, EECBS, with 0.1 suboptimality, outperformed the proposed approaches regarding both solution and optimality rates. Nevertheless, we emphasize that our approach concerns the agents' privacy; therefore, there is a trade-off between privacy and performance. When agents utilize zero or dynamic commitment types, they perform better regarding success rates than when adopting a standard commitment. Furthermore, we observed that the average number of negotiations in each time step and information sharing rates are higher when FoV increases.

In future work, we are planning to extend our approach by adopting multilateral negotiation instead of multiple consecutive bilateral negotiations and to compare its performance with the proposed method. Moreover, it is necessary to design more sophisticated agents capable of thinking ahead in their negotiation and path-planning processes to improve the solution quality and rate. Additionally, when a negotiation fails, our framework stops the run for the entire scenario. However, some of the agents can still continue negotiating with each other to reach their destination, while some agents fail. It would be interesting to examine how many agents can reach their destinations in such settings. Without a doubt, the trade-off between privacy and performance should be further investigated elaborately.

Declarations

Ethical Approval: Not applicable to this study.

Competing interests: The authors declare that they have no conflict of interest.

Authors' contributions: All authors have equal contributions to this study.

Funding: No funding for this study.

Availability of data and materials: The authors can share data and materials upon request.

References

1. Amir, O., Sharon, G., Stern, R.: Multi-agent pathfinding as a combinatorial auction. In: Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI'15, p. 2003–2009. AAAI Press (2015)
2. Aydoğan, R., Hindriks, K.V., Jonker, C.M.: Multilateral Mediated Negotiation Protocols with Feedback, pp. 43–59. Springer Japan, Tokyo (2014)
3. Aydoğan, R., Festen, D., Hindriks, K., Jonker, C.: Alternating offers protocols for multilateral negotiation. *Studies in Computational Intelligence* **674**, 153–167 (2017)
4. Baarslag, T., Gerding, E.H., Aydoğan, R., Schraefel, M.: Optimal negotiation decision functions in time-sensitive domains. In: 2015 IEEE/WIC/ACM International Joint Conferences on Web Intelligence (WI) and Intelligent Agent Technologies (IAT), vol. 2, pp. 190–197 (2015)
5. Bnaya, Z., Stern, R., Felner, A., Zivan, R., Okamoto, S.: Multi-agent path finding for self interested agents. In: Sixth Annual Symposium on Combinatorial Search (2013)
6. Bonisoli, A., Gerevini, A., Saetti, A., Serina, I.: A privacy-preserving model for the multi-agent propositional planning problem. In: ECAI (2014)
7. Brafman, R.I., Domshlak, C.: From one to many: Planning for loosely coupled multi-agent systems. In: ICAPS (2008)
8. Cole, R.J., Dodis, Y., Roughgarden, T.: Pricing network edges for heterogeneous selfish users. In: STOC '03 (2003)
9. Desaraju, V.R., How, J.P.: Decentralized path planning for multi-agent teams with complex constraints. *Auton. Robots* **32**(4), 385–403 (2012). DOI 10.1007/s10514-012-9275-2. URL <https://doi.org/10.1007/s10514-012-9275-2>
10. Eran, C., Keskin, M.O., Cantürk, F., Aydoğan, R.: A decentralized token-based negotiation approach for multi-agent path finding. In: EUMAS (2021)
11. Erdmann, M., Lozano-Perez, T.: On multiple moving objects. In: Proceedings. 1986 IEEE International Conference on Robotics and Automation, vol. 3, pp. 1419–1424 (1986). DOI 10.1109/ROBOT.1986.1087401
12. Felner, A., Stern, R., Shimony, S.E., Boyarski, E., Goldenberg, M., Sharon, G., Sturtevant, N.R., Wagner, G., Surynek, P.: Search-based optimal solvers for the multi-agent pathfinding problem: Summary and challenges. In: SOCS (2017)
13. Gautier, A., Stephens, A., Lacerda, B., Hawes, N., Wooldridge, M.: Negotiated path planning for non-cooperative multi-robot systems. In: Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems, AAMAS '22, p. 472–480. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2022)
14. Ho, F., Geraldès, R., Gonçalves, A., Rigault, B., Sportich, B., Kubo, D., Cavazza, M., Prendinger, H.: Decentralized multi-agent path finding for uav traffic management. *IEEE Transactions on Intelligent Transportation Systems* **23**(2), 997–1008 (2022). DOI 10.1109/TITS.2020.3019397

15. Inotsume, H., Aggarwal, A., Higa, R., Nakadai, S.: Path negotiation for self-interested multirobot vehicles in shared space. In: Proceedings of the International Conference on Intelligent Robots and Systems, pp. 11587–11594. IEEE, Las Vegas, NV, USA (2020)
16. Klein, M., Faratin, P., Sayama, H., Bar-Yam, Y.: Protocols for negotiating complex contracts. *IEEE Intelligent Systems* **18**, 32–38 (2003). DOI 10.1109/MIS.2003.1249167
17. Li, J., Ruml, W., Koenig, S.: Eecbs: A bounded-suboptimal search for multi-agent path finding. In: AAAI (2021)
18. Li, Q., Gama, F., Ribeiro, A., Prorok, A.: Graph neural networks for decentralized multi-robot path planning. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE (2020). DOI 10.1109/iros45743.2020.9341668
19. de Oliveira Ramos, G., Rădulescu, R., Nowé, A., Tavares, A.R.: Toll-based learning for minimising congestion under heterogeneous preferences. In: AAMAS (2020)
20. Parkes, D.C.: Ibundle: An efficient ascending price bundle auction. In: Proceedings of the 1st ACM Conference on Electronic Commerce, EC '99, p. 148–157. Association for Computing Machinery, New York, NY, USA (1999). DOI 10.1145/336992.337032. URL <https://doi.org/10.1145/336992.337032>
21. Patwardhan, A., Murai, R., Davison, A.J.: Distributing collaborative multi-robot planning with gaussian belief propagation. *IEEE Robotics and Automation Letters* **8**(2), 552–559 (2023). DOI 10.1109/LRA.2022.3227858
22. Pritchett, A., Genton, A.: Negotiated decentralized aircraft conflict resolution. *IEEE Transactions on Intelligent Transportation Systems* **19**, 81–91 (2018)
23. Purwin, O., D’Andrea, R., Lee, J.: Theory and implementation of path planning by negotiation for decentralized agents. *Robotics Auton. Syst.* **56**, 422–436 (2008)
24. Rosenschein, J.S., Zlotkin, G.: Rules of Encounter: Designing Conventions for Automated Negotiation among Computers. MIT Press, Cambridge, MA, USA (1994)
25. Roughgarden, T., Tardos, É.: How bad is selfish routing? *J. ACM* **49**, 236–259 (2002)
26. Salzman, O., Stern, R.: Research challenges and opportunities in multi-agent path finding and multi-agent pickup and delivery problems. In: Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS ’20, p. 1711–1715. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC (2020)
27. Sharon, G., Stern, R., Felner, A., Sturtevant, N.R.: Conflict-based search for optimal multi-agent pathfinding. *Artif. Intell.* **219**, 40–66 (2012)
28. Stern, R.: Multi-Agent Path Finding – An Overview, pp. 96–115. Springer International Publishing, Cham (2019)

29. Stern, R., Sturtevant, N., Felner, A., Koenig, S., Ma, H., Walker, T.T., Li, J., Atzmon, D., Cohen, L., Kumar, T.K.S., Boyarski, E., Barták, R.: Multi-agent pathfinding: Definitions, variants, and benchmarks. CoRR **abs/1906.08291** (2019). URL <http://arxiv.org/abs/1906.08291>
30. Sujit, P.B., Sinha, A., Ghose, D.: Multiple uav task allocation using negotiation. In: AAMAS '06: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems, AAMAS '06, p. 471–478. Association for Computing Machinery, New York, NY, USA (2006). DOI 10.1145/1160633.1160719. URL <https://doi.org/10.1145/1160633.1160719>
31. Thayer, J., Ruml, W.: Bounded suboptimal search: A direct approach using inadmissible estimates. In: IJCAI (2011)