

Real-Time Decentralized Navigation of Nonholonomic Agents Using Shifted Yielding Areas

Liang He¹, Zherong Pan², Dinesh Manocha³

Abstract— We present a lightweight, decentralized algorithm for navigating multiple nonholonomic agents through challenging environments with narrow passages. Our key idea is to allow agents to yield to each other in large open areas instead of narrow passages, to increase the success rate of conventional decentralized algorithms. At pre-processing time, our method computes a medial axis for the freespace. A reference trajectory is then computed and projected onto the medial axis for each agent. During run time, when an agent senses other agents moving in the opposite direction, our algorithm uses the medial axis to estimate a Point of Impact (POI) as well as the available area around the POI. If the area around the POI is not large enough for yielding behaviors to be successful, we shift the POI to nearby large areas by modulating the agent’s reference trajectory and traveling speed. We evaluate our method on a row of 4 environments with up to 15 robots, and we find our method incurs a marginal computational overhead of 10-30 ms on average, achieving real-time performance. Afterward, our planned reference trajectories can be tracked using local navigation algorithms to achieve up to a 100% higher success rate over local navigation algorithms alone.

I. INTRODUCTION

In recent years, autonomous vehicles have been deployed in complex, city-scale scenarios to accomplish various tasks such as food delivery, warehouse administration, and public transportation. These vehicles routinely travel on highly regulated paths, such as highways, crossroads, and sidewalks, or in spaces with large open areas including shopping malls, school libraries, etc. Most prior works [1, 2, 3, 4] build navigation algorithms on one of these assumptions. In reality, however, autonomous vehicles must also be prepared for unexpected and unregulated scenarios or spaces with narrow passages. Dealing with narrow spaces is inevitable when two food delivery robot meets in the aisle of a hotel or an autonomous truck travels downtown to reach a warehouse. Narrow passages are notoriously difficult to handle, even when navigating a single robot [5], and scaling to multiple agents is still an open problem.

Prior methods for navigating multiple agents are classified into decentralized local techniques and centralized global techniques, each having its pros and cons. Local navigation methods [1, 2] assume agents move towards their goal positions along some local directions without communicating with each other. When obstacles or other agents get in the way, heuristic behaviors, such as yielding [6], grouping[7], and following [8, 9] are used to avoid collisions. However, local techniques can fail in the face of narrow passages where




Fig. 1: Two agents travel in an environment with one large and one narrow space. (a): Top picture: if the two agents travel at the same speed, their meeting point (defined as the POI and illustrated as the red dot) will be in the narrow space where local navigation techniques can fail. (b)Bottom picture: our method shifts the POI to the large space so that local navigation can successfully generate collision-free trajectories by yielding.

agents form deadlock configurations as illustrated in Figure 1 (a). On the other hand, global navigation methods [10, 11, 12] coordinate agent motions in a central node to avoid collisions. Although these methods can handle many agents in complex environments with narrow passages, they rely on strong assumptions such as the environment being grid-like, agents moving on discrete graph-like structures, or the agents being holonomic. However, actions such as constructing such discrete structures or generalizing to nonholonomic agents are non-trivial and cannot be used in time-critical applications due to a high computational cost.

Main Result: We propose an improved decentralized algorithm for nonholonomic multi-agent navigation, which incorporates ideas from centralized techniques to alleviate the deadlock problem. We observe that yielding behaviors used by prior local navigation approaches [1, 2] can have high success rates in large open areas while being less successful in narrow spaces as illustrated in Figure 1 (b). As a result, we propose shifting the yielding areas to large open spaces of the environment to increase the success rate. Specifically, our algorithm relies on the construction of a medial axis for the free space. By mapping agent positions and their trajectories to the medial axis, we can estimate their Positions-Of-Impacts (POIs), which are positions where agents get close enough for local navigation techniques to generate yielding behaviors. We then estimate the surrounding space required by such yielding behaviors. If the space

¹Liang He is with the University of North Carolina at Chapel Hill. {lianghe.hust@gmail.com} ²Zherong Pan is with the University of North Carolina at Chapel Hill. {zherong.pan.usa@gmail.com} ³Dinesh Manocha is with the University of Maryland, College Park. {dm@cs.umd.edu}

around a POI is not large enough for the yielding to be successful, we search for nearby large spaces and re-plan agent trajectories to move the POI. We show that such re-planning can be accomplished at a relatively low-cost without communication with other agents, preserving the decentralized nature of our method.

We evaluate our method in 4 challenging scenarios with 5-15 robots. The results show that our method exhibits real-time performance, taking up to 20 ms and 43 ms on average to plan the POIs. Compared with local navigation alone, our method achieves up to a 100% higher success rate in some scenarios.

II. RELATED WORK

Over the last two decades, a large body of works on the multi-agent narrow passage navigation problem in motion planning has emerged.

Widely used sampling algorithms such as RRT [13] and PRM [14] can work in high-dimensional configuration spaces by, looking for feasible motion plans, and extensions including RRT* [15] and FMT* [16] can find (nearly) optimal trajectories. These algorithms have been extended to handle nonholonomic agents [17, 18]. Unfortunately, both theoretical analysis [19] and empirical studies [5] have shown that such algorithms incur extremely high computational overheads. Indeed, narrow passages significantly reduce the set of the lookout [20], which is crucial to the efficacy of sampling, while the complexity of optimal motion planning grows exponentially with the number of agents [19]. Almost all these algorithms are offline and inappropriate for time-critical applications such as autonomous driving.

Local navigation techniques use a set of heuristic rules to generate moving directions. These methods incur a much lower computational cost but sacrifice completeness or feasibility. In practice, however, they can have a high success rate under certain assumptions. Successful local navigation algorithms include the dynamic windows [21], reciprocal velocity obstacles (RVO) [6, 1, 2], and potential fields [22, 23]. All these methods were originally proposed for holonomic robots and extensions to differential drive models have been proposed. It is noteworthy that RVO and its variants can provide a collision-free guarantee, which allows agents to alter their moving directions or come to a full stop before collisions. This feature of RVO typically produces a yielding behavior allowing agents to move around local obstacles and continue towards the goal. However, the ambient space required for such yielding behaviors is generally larger for nonholonomic robots than holonomic ones, making RVO-based methods less successful in differential drive models and narrow passages.

A different category of methods, known as centralized, global algorithms [24, 25, 11], involves discretizing the agent motions on a grid or a graph-like structure. Graph search algorithms can then be used to find optimal [26], near optimal [11], or feasible trajectories [27] for large groups of agents within a relatively small computational budget. However, these methods are mostly designed for holonomic

robots, and extensions to nonholonomic cases are far from trivial while their computational cost cannot meet real-time requirements. Our method can be interpreted as a special kind of centralized algorithm on the medial axis graph of the free space, on which we plan the POIs. The low-level yielding actions are then generated using local navigation techniques within each POI.

Finally, we have noticed some recent works [28, 2, 29, 30] apply data-driven techniques to multi-agent navigation problems. By presenting agents with examples of optimal solutions in challenging scenarios, some learned policies can outperform analytic techniques. These techniques are parallel and orthogonal to our contribution. We speculate that learning-based techniques can be used as the local navigator in our method to generate high-quality yielding behaviors in large open areas. However, these methods incur a high computational cost in the training phase and re-training is required when the environment changes. The results of learned navigation policies are also sensitive to training parameters and network architectures. These potential drawbacks inspire us to design low-cost algorithms based on existing local navigation algorithms, with a higher success rate.

III. PROBLEM FORMULATION & BACKGROUND

We assume there are N nonholonomic agents with the configuration of i th agent being $x_i(t)$ at time instance the t . The agent moves in a 2D freespace $\mathcal{F} \subset \mathbb{R}^2$ according to the following differential drive model:

$$\dot{x}_i(t) = f_i(x_i, u_i),$$

where u_i is the control signal. With each agent starting from an initial configuration $x_i(0)$, our goal is to find $u_i(t)$ for $t \in [0, T]$ such that $p(x_i(T))$ is close enough to some goal position g_i , where $p(\bullet)$ is the configuration-to-position mapping function. Given g_i , local navigation algorithms [6, 2] would direct agents via a desired velocity v_i^* and modulate u_i to locally avoid collisions. We build our method on the generalized RVO algorithm denoted as a function:

$$u_i(t) \triangleq \text{GRVO}(v_i^*, x_i(t)).$$

Such modulation typically exhibits yielding behaviors allowing a crowd of agents to move around each other and continue towards their respective goals. However, extra space is required for local yielding to be successful. This property is exploited in prior work [31] to design centralized navigation algorithms for holonomic agents, while nonholonomic agents typically require even larger yielding space. The choice of desired velocity is another key to the success of local navigation. A prominent choice is $v_i^* \triangleq g_i - p(x_i(t))$, which is valid in open areas with small obstacles. For more complex or obstacle-rich environments, a set of reference trajectories must be computed to guide agents across large obstacles.

A. Blum Medial-Axis

Our method makes extensive use of the medial axis of \mathcal{F} to 1) estimate the area required by the yielding behavior and 2) compute reference trajectories. The definition of Blum




Fig. 2: We illustrate the discretized medial axis for the U-shaped environment, where red dots belong to V and black dashed edges belong to E . Each $s_i \in V$ is associated with a circular domain (black circle) with radius defined as $r(s_i)$ (black arrow). A reference trajectory is computed by first projecting $x_i(0)$ (blue) and g_i (green) to G and then computing the shortened path on G .

medial-axis [32] or skeleton is as follows. Given a 2D object defined by a **closed, oriented** boundary $\partial\mathcal{F}$, a Blum medial axis is a set. For every point s in this set, we can find a unique circle centered at s that is tangent to at least two points of $\partial\mathcal{F}$. This circle is known as the circular domain or domain of s and we denote its radius as $r(s)$. A practical method like [33] would compute a discretized Blum medial axis, which is a graph $G = \langle V, E \rangle$, where the set of vertices is sampled skeleton points $V = \{s_i\}$ at regular intervals connected by edges in E . As illustrated in Figure 2, we compute a reference trajectory for the i th agent by first projecting $x_i(0)$ and g_i to the closest vertices and then searching for a trajectory along G via Dijkstra's algorithm.

B. Trajectory Following with Yielding

Given a reference trajectory, we have x_i track the trajectory by designing the desired velocity v_i^* . Specifically, we set the desired velocity to be the negative gradient of a cost function $v_i^* \triangleq -\nabla_{p(x_i)} - c(p(x_i))$ defined as:

$$c(p(x_i)) \triangleq c_{\text{follow}}(p(x_i)) + c_{\text{bias}}(p(x_i)),$$

where c_{follow} guides x_i to move forward along the reference trajectory and c_{bias} penalizes bias from the trajectory. We use an idea similar to the Frenet-frame-based tracking method [34]. Specifically, we first compute the closest $s_i \in V$ to $x_i(t)$ that belongs to the reference trajectory. We denote s_{i+1} as the next node in V that also belongs to the reference trajectory, then we define:

$$c_{\text{follow}}(p(x_i)) \triangleq -(s_{i+1} - s_i)^T \dot{p}(x_i)$$

$$c_{\text{bias}} \triangleq \|p(x_i) - s_i\|^2.$$

In the next section, we describe a method to avoid deadlock configurations in narrow passages, allowing the yielding behaviors generated by GRVO to have a high success rate.

IV. GRVO WITH SHIFTED YIELD AREAS

Our method differs from prior works by applying an additional modulation to the desired velocity function v_i^* and we denote this function as $\mathcal{M}(v_i^*)$. The modulated velocity can be plugged into GRVO to derive our final local navigation algorithm:

$$u_i(t) \triangleq \text{GRVO}(\mathcal{M}(v_i^*), x_i(t)).$$

Note that our method can also be combined with local navigation methods other than GRVO. Our modulation function aims at shifting the POI between the two agents to large open areas in \mathcal{F} . Being a decentralized algorithm, such modulation is highly challenging because an agent does not have the ability to acquire other agents' trajectories, nor to alter their motions. However, we find it suffices to only modulate the velocity of the agent being considered based on a rough estimation of other agents' trajectories, as long as the same modulation function \mathcal{M} is deployed on all the agents. In the following sections, we present details about POI detection, shifting, and modulation.

A. POI Detection




Fig. 3: We illustrate the procedure of estimating the POI between x_i (blue) and x_j (green). Two points are first projected to the closest skeletal nodes s_i and s_j , respectively. The shortened path P_{ij} is illustrated with red nodes and dashed lines. If the difference between $\dot{p}(x_i)$, $\dot{p}(x_j)$ and tangents of P_{ij} are smaller than a user-defined ϵ (gray cones) path, then we assume a POI exists. The POI is the point where two agents meet along P_{ij} . In this figure, since x_j is slower (shorter green arrow), the POI is closer to x_j .

We define for each agent x_i a sensing radius R_i . When any other agent x_j satisfies $\|p(x_i) - p(x_j)\| \leq R_i$, we assume a potential yielding behavior might happen between them. Since x_i does not know x_j 's future trajectory, we need to estimate POI based on the following assumption. We first project $p(x_i), p(x_j)$ onto their closest points on G , which are denoted as s_i and s_j , respectively. We then compute a shortened path between s_i and s_j on G via Dijkstra's algorithm. In practice, we precompute the all-pair shortest distances so any shortened path can be looked up instantaneously. This path is denoted as $P_{ij}(\alpha)$, where $\alpha \in [0, 1]$, $P_{ij}(0) = s_i$, and $P_{ij}(1) = s_j$. If both x_i and x_j are moving along the opposite tangential directions of P_{ij} , then we assume P_{ij} is the estimated path containing a POI of the two agents. We determine that the two agents are traveling along opposite tangential directions if the following conditions hold:

$$\frac{\dot{P}_{ij}(0)^T \dot{p}(x_i)}{\|\dot{P}_{ij}(0)\| \|\dot{p}(x_i)\|} > 1 - \epsilon \quad - \frac{\dot{P}_{ij}(1)^T \dot{p}(x_j)}{\|\dot{P}_{ij}(1)\| \|\dot{p}(x_j)\|} > 1 - \epsilon, \quad (1)$$

and no POI would be considered otherwise. Here ϵ is a user-defined upper bound of velocity bias. For a decentralized algorithm, our agent x_i does not know the velocity of x_j either, so we estimate $\dot{p}(x_j)$ using a finite difference of two

Algorithm 1: POISet(x_i)

```

1: Set ← ∅
2: for Each agent  $x_j \neq x_i \wedge \|p(x_j) - p(x_i)\| < R$  do
3:   Project to skeletal point  $s_i, s_j$ 
4:   Loop up shortest path  $P_{ij}$ 
5:   if Equation 1 holds then
6:     Set ← Set  $\cup \{P_{ij}(\alpha_{POI})\}$ 
7: Return Set

```

consecutive frames of x_j . The POI between x_i and x_j is then estimated as $P_{ij}(\alpha_{POI})$ where α_{POI} is computed such that the following condition holds:

$$\frac{|P_{ij}([0, \alpha_{POI}])|}{\|\dot{p}(x_i)\|} = \frac{|P_{ij}([\alpha_{POI}, 1])|}{\|\dot{p}(x_j)\|},$$

where $|P_{ij}|$ denotes the arc-length of a sub-trajectory. The POI detection procedure is illustrated in Figure 3 and outlined in Algorithm 1, which incurs marginal overhead to conventional local navigation techniques.

B. POI Shifting

Given a POI located at $s_i \in V$, we then estimate its surrounding area. Given the medial axis, this area can be immediately estimated as the circular domain at s_i . If a POI is located on an edge of E neighboring s_i and s_j , we interpolate the circular domain radius. The radius of circular domain $r(s_i)$ must be sufficiently large for the yielding behavior to have a high success rate. Unfortunately, we are still lacking a theoretical analysis connecting the success rate of GRVO and the size of the yielding area. Instead, we use the following heuristic rule to compute the minimal domain radius $r(s_i)$ for n agents to successfully yield to each other:

$$r(s_i) \geq \eta r(n+1), \quad (2)$$

where $\eta \in (0, 1]$ is a user-provided parameter. In typical scenarios, we have $n = 2$ since POI is estimated for two agents. If Equation 2 is violated, we need to shift POI to a nearby large space on the medial axis graph G . We propose first searching for nodes belonging to P_{ij} . This is because P_{ij} lies on our estimated path and shifting POI within P_{ij} would not cause a detour. If P_{ij} does not contain any node satisfying Equation 2, we search the entire G for the nearest node, which is the center of a large domain. If both attempts fail, we decide the entire map consists of narrow spaces and do not shift POI. This procedure is summarized in Algorithm 2.

C. POI Merging

We found that handling only POI cases with two agents improves the success rate of GRVO. For extremely challenging environments, however, more agents can meet at nearby POIs and we must consider POIs involving $n > 2$ agents. We handle this case by iteratively merging nearby POIs as outlined in Algorithm 3. In practice, if two POIs denoted as POI_i and POI_j involve n_i and n_j agents, respectively, we

Algorithm 2: POIShift(POI, n)

```

1: Dist ← ∞,  $POI_0 \leftarrow POI$ ,  $POI \leftarrow \text{None}$ 
2: for Each  $s_i \in P_{ij}$  do
3:   if Equation 2  $\wedge \|POI_0 - p(s_i)\| < \text{Dist}$  then
4:      $POI \leftarrow p(s_i)$ ,  $\text{Dist} \leftarrow \|POI_0 - p(s_i)\|$ 
5: if  $\text{Dist} < \infty$  then
6:   Return  $POI$ 
7: for Each  $s_i \in V$  do
8:   if Equation 2  $\wedge \|POI_0 - p(s_i)\| < \text{Dist}$  then
9:      $POI \leftarrow p(s_i)$ ,  $\text{Dist} \leftarrow \|POI_0 - p(s_i)\|$ 
10: Return  $POI$ 

```

Algorithm 3: POIMerge(x_i)

```

1: Set ← POISet( $x_i$ ), More ← True
2: for  $POI \in \text{Set}$  do
3:   if POIShift( $POI, 2$ )  $\neq \text{None}$  then
4:     Set ← Set / { $POI$ }
5:     Set ← Set  $\cup$  POIShift( $POI, 2$ )
6: while More do
7:   More ← False
8:   for A pair of  $POI_i, POI_j \in \text{Set}$  with  $n_i, n_j$  do
9:     if Equation 3 holds then
10:      if POIShift( $POI_i, n_i + n_j$ )  $\neq \text{None}$  then
11:        Set ← Set / { $POI_i, POI_j$ }
12:        Set ← Set  $\cup$  POIShift( $POI_i, n_i + n_j$ )
13:        More ← True
14:      else if POIShift( $POI_j, n_i + n_j$ )  $\neq \text{None}$  then
15:        Set ← Set / { $POI_i, POI_j$ }
16:        Set ← Set  $\cup$  POIShift( $POI_j, n_i + n_j$ )
17:        More ← True
18: Return Set

```

merge them into a single POI_{ij} if the following condition holds:

$$\|POI_i - POI_j\| \leq \eta \min(r(n_1 + 1), r(n_2 + 1)). \quad (3)$$

The merged POI_{ij} involves $n = n_i + n_j$ agents and its required yielding radius is specified by Equation 2. We perform the POI shifting procedure as described in Section IV-B. If the shifting procedure fails, then we reject merging. We iteratively merge POIs until no more merging can be performed.

D. Velocity Modulation

After the above procedure, an agent has a set of POI positions against a multitude of other agents. We choose the nearest POI to $p(x_i(t))$ as the temporary goal point to modulate our velocity. Note that due to various sources of uncertainty and inaccuracy in estimating POI, modulating our velocity can cause detours. To minimize this effect, we only adopt modulation if the nearest POI was successfully

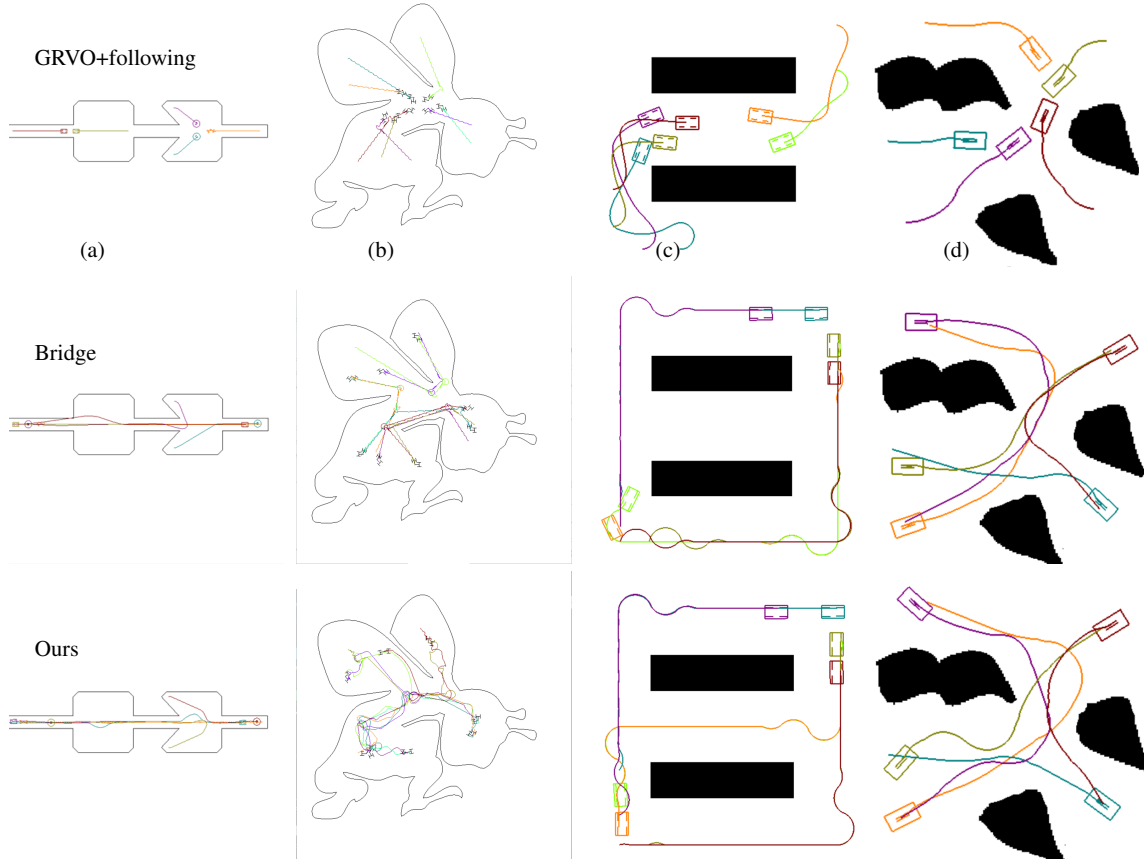


Fig. 4: The 4 most challenging benchmarks used in our experiments with agent trajectories computed using GRVO+following, Bridge, and our method are highlighted in the first, second, and third rows, respectively: (a): We random generate several agents in two open areas connected by a narrow passage; (b): Random agents are generated in a complex Bee-shaped environment; (c): A maze with two long obstacles creating narrow aisles in between; (d): A narrow passage with a garage in the middle. Agents traveling through the narrow passage must move to the garage temporarily to accomplish yielding behaviors.

Method	Benchmark 1			Benchmark 2			Benchmark 3			Benchmark 4		
	Traj. Length	Succ. Rate	FPS	Traj. Length	Succ. Rate	FPS	Traj. Length	Succ. Rate	FPS	Traj. Length	Succ. Rate	FPS
GRVO	1311	0%	45	632	4%	43	433	50%	55	401	50%	60
GRVO+IC	1258	0%	35	562	12%	33	341	94%	42	333	100%	55
Bridge	607	100%	-	424	100%	-	324	100%	-	312	100%	-
Ours	805	100%	41	455	100%	43	421	100%	42	387	100%	43

TABLE I: We compare our approach with previous methods, GRVO, GRVO+following, and Bridge, on the 4 benchmarks, in terms of average agent trajectory length (we only consider trajectories of agents that successfully reach goals), success rate over 50 random scenarios (we only consider a scenario successfully handled when all the agents reach their goals), and the FPS.

shifted, i.e. we define \mathcal{M} as:

$$\mathcal{M}(v_i^*) \triangleq \begin{cases} \left[\underset{\text{POI}}{\operatorname{argmin}} \| \text{POI} - p(x_i) \| \right] - x_i & \text{POI shifted} \\ v_i^* & \text{otherwise.} \end{cases}$$

E. Acceleration by Precomputation

The main computational bottleneck of our algorithm lies in the POI merging procedure. This involves at most $2N$ calls to Algorithm 2, and each call to Algorithm 2 incurs a computational cost of $|V| + |E|$. However, we can further reduce the cost of Algorithm 2 to $\mathcal{O}(1)$ by precomputing a lookup table. Note that POI generally lies on an edge of G . However, if we use sufficiently dense samples to construct V , we can shift POI to a nearby vertex $s_i \in V$ incurring a small error. In this way, Algorithm 2 will only be called

with discrete inputs $\text{POIShift}(s_i, n)$, and we can construct a table of size $|V| \times N$ to precompute all possible results. After such acceleration, the complexity of each evaluation of modulation function \mathcal{M} is only $\mathcal{O}(2N)$.

V. EVALUATION

We have implemented our algorithms in C++ on an Intel Core i7 CPU running with 16GB of RAM. We use the CGAL library to build the medial axis graph G . We evaluate our method on three categories of robots: the single differential-drive robot, the Dubin's car, and the differential-drive robot with trailer (truck for short) as in [2], where we tune the parameters such that the maximal turning curvature of the trajectory is 0.19 for a Dubin's car, and 0.22 for a truck-like robot. In all three testing scenarios, we use GRVO [2]

as our local navigation algorithm. We set the vehicle size to 4×5 square units and, the medial axis sampling interval to 0.02 units, and we use $\eta = 1.6 - 2.6$, $R = 15 - 30$. We randomly put the agents in the open area and repeat 50 times in each scenario. The average computational cost over these scenarios is summarized in Figure 5, and it largely depends on the number of robots. Our bottleneck lies in the collision detection between robots of non-circular shapes.




Fig. 5: We plot the averaged computational cost of our method (frames per second) against the number of robots.

A. Baselines

There are several prior works on improving the success rate of local navigation methods. Our first baseline is the GRVO algorithm [2] without our modulation. Our second baseline is the grouped sampling-based algorithm (Bridge) [35]. This algorithm aims at solving the same problem as ours.

They use a sampling-based method to precompute a set of corridors across narrow passages, in which nonholonomic agent trajectories can be efficiently generated by interpolation. Agents follow these interpolated trajectories in the corridor while collisions are handled using local navigation techniques. Finally, we also consider the GRVO algorithm with adapted Implicit Coordination method (GRVO+following) [9]. As the major difference from our method, the IC algorithm allows an agent to communicate with neighbors to coordinate desired velocities.

B. Benchmark Problems

We consider 4 challenging benchmarks, illustrated in Figure 4. The trajectories generated by different methods are compared and evaluated in three aspects: the average length of agent trajectories, the rate of success of finding feasible motion plans, and the frame rate per second (FPS). Quantitative results, corresponding to an average of over 50 simulations with randomly generated agent configurations in an assigned sub-area of the scenarios, are summarized in Table I.

	Bridge	Ours
I	212	0.6
II	122	3.3
III	12	0.2
IV	19	4

TABLE II: This table shows the average computational time (seconds) during the precomputation stage of Bridge and our method on 4 scenarios.

Benchmark I: We use a dumb-like environment, shown in Figure 4 (a), where multiple agents move from one side to the other. We observe that our method allows the agents to determine that they will meet agents from the other side and the POIs lie in the narrow central passage. Our method then has agents on one side retreat from the narrow space and return to the left side of the scene to wait for agents from the other side to pass through before moving on. For this example, most other local navigation methods, including GRVO and GRVO+following, fail.

Benchmark II: As shown in Figure 4 (b), we place a group of agents in a complex, bee-shaped environment. Agents start from one corner of the freespace and repeatedly yield other upcoming agents. The results in Table I show that our method can always compute a feasible motion plan, while prior techniques cause many collisions resulting in deadlock configurations. The only rival algorithm that exhibits a high success rate is Bridge, which uses a sampling-based motion planner during the precomputation stage. In comparison, the precomputation involved in our method is only used to find the medial axis graph of \mathcal{F} , which can be accomplished much faster as profiled in Table II.

Benchmark III: As shown in Figure 4 (c), we use a small maze involving two long obstacles, with agents again placed randomly. The Bridge algorithm outperforms our method for this benchmark in terms of trajectory length, although the success rates of both algorithms are 100%. This is due to inaccuracies in detecting and shifting POIs, where our method does not account for in-between obstacles.

Benchmark IV: Our last and most challenging benchmark involves a single narrow passage with a garage in the middle for agents to perform yielding. As illustrated in Figure 4 (d), our POI shifting procedure allows agent to be directed to the garage, while all prior methods fail.

VI. CONCLUSION & LIMITATION

We propose a novel velocity modulation algorithm to improve the success rate of prior local navigation algorithms for multiple nonholonomic agents. We observe that local navigation methods can generate yielding behaviors for agents so they can move around each other and continue toward their respective goals. However, yielding requires extra space and can have a low success rate in narrow passages. To alleviate this problem, we propose shifting the POI between two or more agents to large open areas. We show that even using a rough estimation of POI and the required space for yielding, such a strategy can empirically improve the success rate of conventional local navigation algorithms such as GRVO [2] by 100% in some scenarios. A major issue with the current method is our decentralized setting, which does not allow any communication or coordination between agents. It is thus difficult to further improve the accuracy of POI estimation and shifting. In further works, we are considering extending our method to allow local communications between agents to achieve partial coordination, as has been done in [36].

REFERENCES

- [1] J. Alonso-Mora, A. Breitenmoser, P. Beardsley, and R. Siegwart, "Reciprocal collision avoidance for multiple car-like robots," in *2012 IEEE International Conference on Robotics and Automation*, 2012, pp. 360–366.
- [2] D. Bareiss and J. van den Berg, "Generalized reciprocal collision avoidance," *The International Journal of Robotics Research*, vol. 34, no. 12, pp. 1501–1514, 2015.
- [3] M. Whitzer, D. Shishika, D. Thakur, V. Kumar, and A. Prorok, "Dc-capt: Concurrent assignment and planning of trajectories for dubins cars," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 8791–8797.
- [4] L. Claussmann, M. Revilloud, D. Gruyer, and S. Glaser, "A review of motion planning for highway autonomous driving," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 5, pp. 1826–1848, 2020.
- [5] J. Szkandera, I. Kolingerová, and M. Maňák, "Narrow passage problem solution for motion planning," in *International Conference on Computational Science*, Springer, 2020, pp. 459–470.
- [6] J. Van Den Berg, S. J. Guy, M. Lin, and D. Manocha, "Reciprocal n-body collision avoidance," in *Robotics research*, Springer, 2011, pp. 3–19.
- [7] L. He and J. van den Berg, "Meso-scale planning for multi-agent navigation," in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 2839–2844.
- [8] L. He, J. Pan, S. Narang, and D. Manocha, "Dynamic group behaviors for interactive crowd simulation," in *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, ser. SCA '16, Zurich, Switzerland: Eurographics Association, 2016, 139–147.
- [9] L. He, J. Pan, W. Wang, and D. Manocha, "Proxemic group behaviors using reciprocal multi-agent navigation," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 292–297.
- [10] J. van den Berg, J. Snoeyink, M. Lin, and D. Manocha, "Centralized path planning for multiple robots: Optimal decoupling into sequential plans," in *Proceedings of Robotics: Science and Systems*, Seattle, USA, 2009.
- [11] J. Yu and D. Rus, "An effective algorithmic framework for near optimal multi-robot path planning," in *Robotics research*, Springer, 2018, pp. 495–511.
- [12] L. He, Z. Pan, K. Solovey, B. Jia, and D. Manocha, "Multi-robot path planning using medial-axis-based pebble-graph embedding," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, IEEE, 2022, pp. 9987–9994.
- [13] S. M. LaValle and J. J. Kuffner Jr, "Randomized kinodynamic planning," *The international journal of robotics research*, vol. 20, no. 5, pp. 378–400, 2001.
- [14] L. Kavraki, P. Svestka, and M. H. Overmars, *Probabilistic roadmaps for path planning in high-dimensional configuration spaces*. Unknown Publisher, 1994, vol. 1994.
- [15] S. Karaman, M. R. Walter, A. Perez, E. Frazzoli, and S. Teller, "Anytime motion planning using the rrt," in *2011 IEEE International Conference on Robotics and Automation*, IEEE, 2011, pp. 1478–1483.
- [16] L. Janson, E. Schmerling, A. Clark, and M. Pavone, "Fast marching tree: A fast marching sampling-based method for optimal motion planning in many dimensions," *The International journal of robotics research*, vol. 34, no. 7, pp. 883–921, 2015.
- [17] D. J. Webb and J. Van Den Berg, "Kinodynamic rrt*: Asymptotically optimal motion planning for robots with linear dynamics," in *2013 IEEE International Conference on Robotics and Automation*, IEEE, 2013, pp. 5054–5061.
- [18] Y. Li, Z. Littlefield, and K. E. Bekris, "Asymptotically optimal sampling-based kinodynamic planning," *The International Journal of Robotics Research*, vol. 35, no. 5, pp. 528–564, 2016.
- [19] L. Janson, B. Ichter, and M. Pavone, "Deterministic sampling-based motion planning: Optimality, complexity, and performance," *The International Journal of Robotics Research*, vol. 37, no. 1, pp. 46–61, 2018.
- [20] D. Hsu, J.-C. Latombe, and R. Motwani, "Path planning in expansive configuration spaces," in *Proceedings of International Conference on Robotics and Automation*, vol. 3, 1997, 2719–2726 vol.3.
- [21] D. Fox, W. Burgard, and S. Thrun, "The dynamic window approach to collision avoidance," *IEEE Robotics & Automation Magazine*, vol. 4, no. 1, pp. 23–33, 1997.
- [22] Y. Koren and J. Borenstein, "Potential field methods and their inherent limitations for mobile robot navigation," in *Proceedings. 1991 IEEE International Conference on Robotics and Automation*, IEEE, 1991, pp. 1398–1404.
- [23] Y. Ma, G. Zheng, W. Perruquetti, and Z. Qiu, "Motion planning for non-holonomic mobile robots using the i-pid controller and potential field," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014, pp. 3618–3623.
- [24] J. Yu and S. M. LaValle, "Multi-agent path planning and network flow," in *Algorithmic foundations of robotics X*, Springer, 2013, pp. 157–173.
- [25] R. J. Luna and K. E. Bekris, "Push and swap: Fast cooperative path-finding with completeness guarantees," in *Twenty-Second International Joint Conference on Artificial Intelligence*, 2011.
- [26] G. Sharon, R. Stern, A. Felner, and N. R. Sturtevant, "Conflict-based search for optimal multi-agent pathfinding," *Artificial Intelligence*, vol. 219, pp. 40–66, 2015.
- [27] J. Yu and D. Rus, "Pebble motion on graphs with rotations: Efficient feasibility tests and planning algorithms," in *Algorithmic foundations of robotics XI*, Springer, 2015, pp. 729–746.
- [28] W. Sun, N. Sood, D. Dey, G. Ranade, S. Prakash, and A. Kapoor, "No-regret replanning under uncertainty," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2017, pp. 6420–6427.
- [29] S. Barrett, P. Stone, S. Kraus, and A. Rosenfeld, "Teamwork with limited knowledge of teammates," in *Twenty-Seventh AAAI Conference on Artificial Intelligence*, 2013.
- [30] P. Trautman, J. Ma, R. M. Murray, and A. Krause, "Robot navigation in dense human crowds: Statistical models and experimental studies of human–robot cooperation," *The International Journal of Robotics Research*, vol. 34, no. 3, pp. 335–356, 2015.
- [31] K. Solovey, J. Yu, O. Zamir, and D. Halperin, "Motion planning for unlabeled discs with optimality guarantees," in *Proceedings of Robotics: Science and Systems*, Rome, Italy, 2015.
- [32] H. Blum and R. N. Nagel, "Shape description using weighted symmetric axis features," *Pattern recognition*, vol. 10, no. 3, pp. 167–180, 1978.
- [33] A. Telea, T. Preusser, C. Garbe, M. Droske, and M. Rumpf, "A variational approach to joint denoising, edge detection and motion estimation," in *Joint Pattern Recognition Symposium*, Springer, 2006, pp. 525–535.
- [34] M. Werling, J. Ziegler, S. Kammel, and S. Thrun, "Optimal trajectory generation for dynamic street scenarios in a frenet frame," in *2010 IEEE International Conference on Robotics and Automation*, IEEE, 2010, pp. 987–993.
- [35] L. He, J. Pan, and D. Manocha, "Efficient multi-agent global navigation using interpolating bridges," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2017, pp. 4391–4398.
- [36] D. Hildreth and S. J. Guy, "Coordinating multi-agent navigation by learning communication," *Proceedings of the ACM on Computer Graphics and Interactive Techniques*, vol. 2, no. 2, pp. 1–17, 2019.