

Probabilistic verification of a decentralized policy for conflict resolution in multi-agent systems

Lucia Pallottino*, Vincenzo Giovanni Scordio*, Emilio Frazzoli[†] and Antonio Bicchi*

* Interdepartmental Research Center “E. Piaggio”, University of Pisa, 56126 Pisa, Italy.

[†] University of California at Los Angeles, Los Angeles, CA, 90095.

Abstract—In this paper, we consider a decentralized cooperative control policy proposed recently for steering multiple non-holonomic vehicles between assigned start and goal configurations while avoiding collisions. The policy is known to ensure safety (i.e., collision avoidance) for an arbitrarily large number of vehicles, if initial configurations satisfy certain conditions. The method is highly scalable, and effective solutions can be obtained for several tens of autonomous agents. On the other hand, the liveness properties of the policy, i.e. the capability of negotiating a solution in finite time, are not completely understood yet. In this paper, we introduce a condition on the final vehicle configurations, which we conjecture to be necessary and sufficient for guaranteeing liveness. We prove the necessity by a constructive method. Because of the overwhelming complexity of proving the sufficiency of such condition, we assess the correctness of the conjecture in probability through the analysis of the results of a large number of randomized experiments.

I. INTRODUCTION

In this paper, we consider the problem of collision-free motion planning for a number of mobile agents evolving on the plane. Agents are modeled as nonholonomic vehicles, constrained to move at constant speed and with bounds on the curvature: such a model for the agent dynamics is very similar to the well-known model for car-like vehicles due to Dubin [1], except that in our case the agents cannot stop but at their targets. The environment in which the agents move is considered to be unbounded and free of obstacles. The agents are aware of the position and orientation of nearby agents, within a certain sensing or communication radius, but have access to no other information. In particular, agents are not supposed to communicate explicitly their goals or their velocities. All agents make decisions based on a common set of rules that are decided a priori, and rely on the assumption that other agents apply the same rules. Some areas of application of the considered problem include air traffic control, manufacturing plants, automated factories, and intelligent transportation systems.

In recent years, the problem of safely coordinating the motion of several robots sharing the same environment has received a great deal of attention, both in robotics and in other application domains. A number of techniques have been developed for omni-directional (holonomic) robots, most of them requiring some form of central authority, either prioritizing robots off-line, or providing an online conflict-resolution mechanism, e.g., [2]–[4]; a characterization of Pareto-optimal solutions has been provided in [5].

Decentralized control policies, acting solely on locally available information, are attractive because of their scalability

to large-scale systems, and of their robustness to single-point failures. However, since the agents act only on local information, global properties of a decentralized control policy are often hard to establish. Several decentralized algorithms have appeared, e.g. [6], [7] for holonomic robots, and [8] for aircraft-like vehicles. The literature on flocking and formation flight, which has flourished recently (e.g., [9]–[11]), while ultimately leading to conflict-free collective motion, does not address individual objectives, and agents are not guaranteed to reach a pre-assigned individual destination. Very recently, Kyriakopoulos and coworkers introduced decentralized control policies ensuring the safe coordination of non-holonomic vehicles [12]. However, the control laws in [12] are not directly applicable to our case, in which vehicles are constrained to move at constant speed, and cannot stop or back up.

In the literature dealing more specifically with air traffic control, the early work of [13] introduced the so-called roundabout technique, which shares some of the qualitative characteristics of the solution considered here. This policy was proven safe for two- and three-aircraft conflicts [14], [15]. A different approach, relying on the solution of Mixed-Integer Linear Programs (MILPs), and on the local exchange of information among “teams” of aircraft, was proven safe (i.e., collision-free) for encounters of up to five aircraft [16]. Remarkably, to the authors’ best knowledge, papers in multi-agent traffic management appear to focus uniquely on proving safety of proposed policies, while the liveness issue (i.e., conflict negotiation in finite time) is typically disregarded.

In this paper, we discuss a control policy, first introduced in [17] which is (i) spatially decentralized, and (ii) provably safe, regardless of the number of vehicles present in the environment. The method builds on [6], wherein the case of holonomic robots moving in an environment with stationary obstacles was considered by introducing a spatially decentralized cooperative control scheme guaranteeing that no collisions occur between robots using limited sensing range.

Although our policy was proven to be safe for an arbitrarily large number of agents, and indeed very effective in negotiating conflicts of several tens of agents, its liveness properties are not yet completely understood. In other words, while it is known that the proposed policy never causes collisions under some mild assumptions on the initial conditions, it was not clear under what conditions on the initial and final configurations the policy ensures that each vehicle will reach

the intended destination in finite time. Notice that in [6], no liveness guarantees were given, and indeed counterexamples were provided.

In considering liveness and safety of the proposed policy, we provide conditions under which both properties are satisfied. Unfortunately, the formal verification that such conditions are sufficient to ensure liveness appears to be overwhelmingly complex. We therefore assess the correctness of the conjecture in probability through the analysis of the results of a large number of randomized experiments.

The study of probabilistic methods for analysis and design of control systems has recently received a growing interest in the scientific community. In particular, probabilistic methods are widely used in robust control [18]. These methods build on the classical Monte Carlo approach and provide theoretically sound justification of results based on probabilistic inequalities theory. Unlike classical worst-case methods, such algorithms provide a probabilistic assessment on the satisfaction of design specifications.

II. BACKGROUND AND PROBLEM FORMULATION

Let us consider n mobile agents moving on the plane at constant speed, along paths with bounded curvature. Let the configuration of the i -th agent be specified by $g_i \in SE(2)$, the group of rigid body transformation on the plane. In coordinates, the configuration of the i -th agent is given by the triple $g_i = (x_i, y_i, \theta_i)$, where x_i and y_i specify the coordinates of a reference point on the agent's body with respect to an orthogonal fixed reference frame, and the heading θ_i is the angle formed by a longitudinal axis on the agent's body with the $y = 0$ axis.

Each agent enters the environment at the initial configuration $g_i(0) = g_{0,i} \in SE(2)$, and is assigned a target configuration $g_{f,i} \in SE(2)$. The agents move along a continuous path $g_i : \mathbb{R} \rightarrow SE(2)$ according to the model

$$\begin{aligned} \dot{x}_i(t) &= v_i \cos(\theta_i(t)) \\ \dot{y}_i(t) &= v_i \sin(\theta_i(t)) \\ \dot{\theta}_i(t) &= \omega_i(t) \end{aligned} \quad (1)$$

where $\omega_i : \mathbb{R} \rightarrow [-\frac{1}{R_C}, \frac{1}{R_C}]$ is a bounded signed curvature control signal. Without loss of generality we can scale the control $\omega_i \in [-1, 1]$ by considering $R_C = 1$. Linear velocity v_i is constant and can be supposed equal to 1 for each agent without loss of generality.

A *collision* is said to occur at time t_c between two agents, if the agents are closer than a specified safety Euclidean distance d_s . Hence, associating to each agent a *safety disc* of radius $R_S = d_s/2$ centered at the agent position a collision occurs whenever two safety discs overlap.

A dynamic feedback control policy π is a map that associates to an individual agent a control input, based on a set of locally-available *internal variables*, and on the current configuration of other agents in the environment. The policy π is said *spatially decentralized* if it is a function only of the configurations of agents that are within a given alert distance d_a from the computing agent.




Fig. 1. The reserved disc of a nonholonomic vehicle with bounded angular velocity.

We describe below a spatially decentralized cooperative policy for collision avoidance, referred to as *Generalized Roundabout Policy*, introduced by the authors in [17].

Reserved disc: The proposed policy is based on the concept of *reserved disc*, over which each active agent claims exclusive ownership. Given the agent configuration g , the associated reserved disc has radius $R = 1 + R_S$, is centered in $(x^c, y^c) = (x + \sin(\theta), y - \cos(\theta))$ and inherits the agent's heading θ (see fig. 1). The configuration $g^c = (x^c, y^c, \theta^c)$ of the reserved disc has the dynamics $\dot{g}^c = ((1 + \omega) \cos \theta^c, (1 + \omega) \sin \theta^c, \omega)$.

Notice that when the agent has control $\omega_i = -1$, corresponding to a maximum curvature radius clockwise turn, the center of the associated reserved disc is fixed, see Figure 1. Hence the reserved disc can be stopped at any time, by setting $\omega = -1$ and it can be moved in any direction, provided one waits long enough for the heading θ to reach the appropriate value.

Constraints: A sufficient condition to ensure safety is that the interiors of reserved disks are disjoint at all times since they always contain the agent's safety discs. If the reserved disk of agent i is in contact with the reserved disks of agents with indices in $J_i \subset \{1, \dots, n\}$, the motion of the agents is constrained as follows

$$\dot{x}_i^c(x_i^c - x_j^c) + \dot{y}_i^c(y_i^c - y_j^c) \geq 0, \quad \forall j \in J_i. \quad (2)$$

In other words, the velocity of the i -th reserved disk is constrained to remain in the convex cone Θ , namely *admissible cone*, determined by the intersection of a number of closed half-planes (2). In the following, we denote with Θ^- the open set obtained removing the boundary of Θ in the clockwise direction. Note that Θ can be computed assuming that each agent is aware of the configuration of all agents within an alert distance $d_a = 4 + d_s$. Hence, the amount of information needed by each agent to compute Θ is bounded and independent from the number of agents in the system: in fact, at each instant, the maximum number of agents with distance less than d_a from the considered agent is six (see fig. 2).

Holding: As previously mentioned, setting $\omega = -1$ causes an immediate stop of an agent's reserved disk's motion. We will say that when $\omega = -1$, the agent is in the *hold* state.

Moving in a free space: In an obstacle free environment, an agent can accomplish the task of reaching an assigned final configuration g_f , starting from g_0 , switching between the




Fig. 2. Maximum number of agents in contact is six.

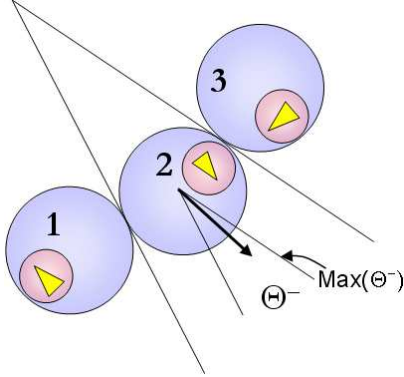


Fig. 3. The set of allowable directions in which the center of the i -th reserved disk can move generated by the contact with reserved discs of vehicles j , m and k respectively.

hold state and the *straight* state associated to the control $\omega = 0$ of the agent. The switching policy can be summarized as follows. Let Δ_f be the vector from the center of the reserved disc g_c to the center g_{cf} of the reserved disc associated with the final configuration. Furthermore, let $\phi : \mathbb{R}^2 \setminus 0 \rightarrow S^1$ be a function returning the polar angle of a vector. Whenever the heading θ of the agent is equal to $\phi(\Delta_f)$ the agent switches its control to $\omega = 0$ and moves straight toward the final configuration until $g_c \equiv g_{cf}$. At this point the agent switches its control to $\omega = -1$ until the target configuration is reached. **Avoiding collision:** As already mentioned in the properties of the reserved region motions, by switching in the *hold* mode, the reserved region stops. Hence, each agent can switch to this mode whenever its heading does not belong to the admissible cone generated by possible contacts between reserved discs, i.e. $\theta \notin \Theta^-$.

Stationary obstacle: If the path of the reserved disk to its position at the target is blocked by another reserved disk, a possible course of action is represented by rolling in a pre-specified direction (in our case, the *positive* direction) on the boundary of the blocking disk. In order to roll on such disk, without violating safety constraints, the control input must be set to $\omega = (1 + R_S)^{-1}$ as soon as the heading of the agent is equal to the value of the counterclockwise direction boundary of the admissible cone, namely $\theta = \max(\Theta^-)$. We refer to this mode as the *roll* state. While the above condition on θ is not true the agent remains in the *hold* state (i.e. $\omega = 0$).

Moving obstacle: In general, the reserved disk of an agent will

not necessarily remain stationary while an agent is rolling on it. While it can be recognized that the interiors of the reserved disks of two or more agents executing the described maneuver will always remain disjoint, it is possible that contact between two agents is lost unexpectedly. In this case, we introduce a new state, which we call *roll2*, in which the agents turns in the positive direction at the maximum rate, i.e., $\omega = +1$, unless this violates the constraints. The rationale for such a behavior is to attempt to recover contact with the former neighbor, and to exploit the maximum turn rate when possible. The *roll2* state can only be entered if the previous state was *roll*. The agent is forced to exit from the *roll2* state after at most time 2π .

Generalized Roundabout Policy: We are now ready to state our policy for cooperative, decentralized, conflict resolution; we call it Generalized Roundabout (GR) policy. The policy followed by each vehicle is based on four distinct *modes of operation*, each assigning a constant value to the control input ω . As a consequence, the closed-loop behavior of an individual agent can be modeled as a hybrid system. We refer the reader to the relevant literature for a more in-depth discussion of the hybrid systems formalism (e.g., [19]–[22] and references therein).

The states of the hybrid systems are 4 ($Q = \{\text{roll}, \text{roll2}, \text{hold}, \text{straight}\}$) and correspond to constant inputs $\omega_{\text{roll}} = (1 + R_S)^{-1}$, $\omega_{\text{roll2}} = +1$, $\omega_{\text{hold}} = -1$, and $\omega_{\text{straight}} = 0$, respectively.

The map Φ_{GR} that describes the agents' dynamic in each node of the system, is derived from (1), substituting the appropriate value for ω , based on the discrete mode, and by the clock rate $\dot{\tau} = 1$ (needed only in the *roll2* state), i.e., it can be written in coordinates as follows:

$$\begin{aligned} \dot{x} &= \cos(\theta) \\ \dot{y} &= \sin(\theta) \\ \dot{\theta} &= \omega_q, \quad q \in Q \\ \dot{\tau} &= 1. \end{aligned} \tag{3}$$

We do not explicitly write down the GR policy and its transition relations, guards, and invariants, but we refer the reader to Figure 4, which should provide the necessary detail in a clearer fashion.

The multiple-vehicle system (\mathcal{S}_{GR}) we are considering is the parallel composition of n agents of the hybrid system described above. We do not define the operation of parallel composition here; see, e.g., [23] for details.

III. ANALYSIS OF THE POLICY

The policy described in the previous section can be shown to provide effective solutions for large-scale problems, such as e.g. the 70-agents conflict resolution illustrated in fig. 5. In this section, we investigate methods to systematically assess conditions under which the policy is applicable and provides solutions which are guaranteed to be collision-free (i.e. *safe* and to ultimately lead all agents to their goals avoiding stalls (i.e. *non-blocking*, or *live*).

Consider a framework in which new agents may issue a request to enter the scenario at an arbitrary time and with

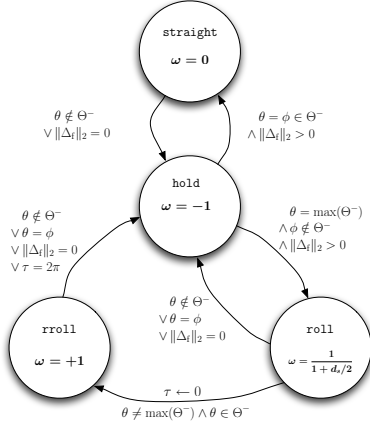


Fig. 4. A hybrid automaton describing the Generalized Roundabout policy.




Fig. 5. A conflict resolution problem with 70 agents in narrow space, for which the proposed policy provides a correct solution. Initial configurations are identified by the presence of gray circles, indicating their reserved discs.

an arbitrary “flight plan”, consisting of an initial and final configuration. In this case, it is important to have conditions to efficiently decide on the acceptability of a new request, i.e. whether the new proposed plan is compatible with safety and liveness of the overall system. The decision whether a new flight plan is admissible may be made by a centralized decision maker, based only on information on the current and final configurations of all agents (real-time collision avoidance remains strictly decentralized, however).

The problem of certifying the admissibility of a requested plan can be dealt with most effectively by decoupling the safety and liveness aspects of current and final configurations. Indeed, for a given policy π , consider the two properties:

P₁: A configuration set $G = \{g_i, i = 1, \dots, n\}$, is *unsafe* for the policy π if there exists a set of target configurations $G_f = \{g_{f,i}, i = 1, \dots, n\}$ such that application of π leads to a collision;

P₂: A target configuration set $G_f = \{g_{f,i}, i = 1, \dots, n\}$, is *blocking* for the policy π if there exists a set of configurations $G = \{g_i, i = 1, \dots, n\}$ from which the application of π

leads to a dead- or live-lock.

A plan $(G(t), G_f)$ is *admissible* if it verifies the predicate $\neg \mathbf{P}_1(G(t)) \wedge \neg \mathbf{P}_2(G_f)$. A simple test to check the first property is provided by the following proposition.

Proposition 1: Property $\mathbf{P}_1(\mathbf{G})$ is verified for the GR policy if and only if the reserved disks of at least two agents in G overlap.

Proof: The proof is a straightforward extension of the safety results provided in [17], and is omitted for brevity. ■

The analysis of property \mathbf{P}_2 is more complex, and hinges upon the definition of a condition concerning the separation of reserved discs associated with target configurations. Let $G_f^c = \{g_{f,i}^c, i = 1, \dots, n\}$ denote the set of configurations of the reserved discs corresponding to G_f , and $P_f^c = \{(x_{f,i}^c, y_{f,i}^c), i = 1, \dots, n\}$ be the set of their center coordinates.

Sparsity condition: for all $(x, y) \in \mathbb{R}^2$ and for $m = 2, \dots, n$,

$$\text{card}\{(x_{f,i}^c, y_{f,i}^c) \in P_f^c : \|(x_{f,i}^c, y_{f,i}^c) - (x, y)\|_2 < \rho(m)\} < m, \quad (4)$$

where

$$\rho(m) = \begin{cases} 2(1 + R_S) & \text{for } m \leq 4, \\ (1 + \cot(\frac{\pi}{m}))(1 + R_S) & \text{for } m \geq 4. \end{cases} \quad (5)$$

In other words, any circle of radius $\rho(m)$, with $1 < m \leq n$, can contain at most $m - 1$ reserved disk centers of targets.

P₃: A target configuration set $G_f = \{g_{f,i}, i = 1, \dots, n\}$, is *clustered* if it violates the sparsity condition.

Proposition 2: Property $\mathbf{P}_2(G_f)$ is verified for the GR policy if $\mathbf{P}_2(G_f)$ is verified, e.g. G_f violates (4).

Proof: The proof is obtained constructively by showing that for all non-sparse target configurations there exists at least an initial condition that, under the GR policy, produces a livelock. Let $\hat{m} \geq 2$ denote the maximum cardinality of subsets of P_f^c that violate the sparsity condition (4), and let $P_{f,\hat{m}}^c \subset P_f^c$ denote one such subset. Take initial conditions for the $n - \hat{m}$ agents corresponding to $P_f^c \setminus P_{f,\hat{m}}^c$ to coincide with their respective targets.

Case $\hat{m} \geq 5$. Consider the smallest circle containing $P_{f,\hat{m}}^c$ and the concentric circle $C_{\hat{m}}$ of radius $\rho(\hat{m}) - (1 + R_S)$. Take initial conditions for the \hat{m} agents such that their reserved discs are centered on $C_{\hat{m}}$ and head in the tangent direction (see fig. 6-a). By applying the GR policy to this configuration, the \hat{m} agents start and stay in *hold* mode until they all reach $\theta_i = \max(\Theta_i^-)$ and switch to the *roll* state. Immediately after the switch, contact between agents is lost, and all switch to *roll12* (fig. 6-b) until contact is re-established, and all switch simultaneously back to *hold*. At this time, agents are in the initial configuration rotated by $2\pi/\hat{m}$ (fig. 6-c). A livelock cycle is thus obtained after \hat{m} such sequences.

Case $2 < \hat{m} \leq 4$. The construction is analogous to the previous case, but $C_{\hat{m}}$ has now radius $\rho(\hat{m})$. Take initial conditions for $\hat{m} - 1$ agents so that their reserved discs are centered on $C_{\hat{m}}$ $2\pi/(\hat{m} - 1)$ radians apart and head in the

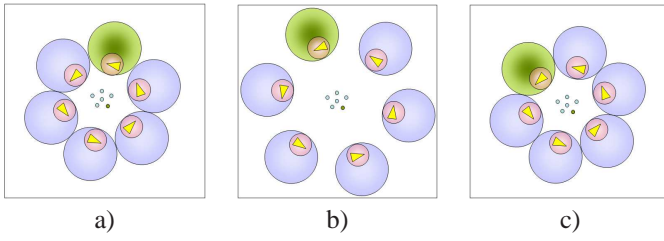


Fig. 6. Livelock-generating conditions for the GR policy with $\widehat{m} = 6$.




Fig. 7. Blocking executions of the GR policy with $\widehat{m} \leq 4$.

tangent direction (see fig. 7-a and -b). Place the initial position of the reserved disc of the remaining agent in the center of $C_{\widehat{m}}$. By applying the GR policy, this agent remains indefinitely in the `hold` state while the other $\widehat{m} - 1$ remain in the `roll` state. Indeed, while in `roll`, the admissible cone coincides with the half plane determined by the tangent to the reserved disc of the inner agent, hence $\theta_i \equiv \max\{\Theta_i^-\} \in \Theta_i^-$. Moreover, by the same reason, $\phi_i \notin \Theta_i^-$. Therefore, no guard leaving `roll` is ever active for these agents.

Case $\widehat{m} = 2$. The construction and behaviour in this case is completely analogous to the case $\widehat{m} \geq 5$ (see fig. 7-c). ■

We have thus proved that sparsity of target configurations is a necessary condition to rule out the possibility of blocking executions of the GR policy. A proof of sufficiency appears to be very complex. In the next section, we describe a method to approach the problem from a probabilistic point of view.

IV. PROBABILISTIC VERIFICATION OF THE GR POLICY

Consider the following statement:

Conjecture The GR policy provides a non-blocking solution for all plans (G_0, G_f) that are safe and verify the sparsity condition.

Let the predicate $\mathbf{P}_{GR}(G_0, G_f)$ be true if the generalized roundabout policy provides a non-blocking solution for initial and final configurations G_0 and G_f , respectively.

A probabilistic verification of the conjecture can be obtained following the approach described below (for more details, see e.g. [18]).

Consider a bounded set $\mathcal{B} = \mathcal{B}_0 \times \mathcal{B}_f$ where the uncertainty $\Delta = (G_0, G_f)$ is uniformly distributed. Let $\mathcal{G} = \{(G_0, G_f) \in \mathcal{B} \mid \mathbf{P}_{GR}(G_0, G_f)\}$ denote the “good” set of problem data for which the predicate applies. Also, let $\mathcal{C} = \{(G_0, G_f) \in \mathcal{B} \mid \neg \mathbf{P}_1(G_0) \wedge \neg \mathbf{P}_3(G_f)\}$ denote the set of safe plans that verify the sparsity condition.

Using the standard induced measure on \mathcal{B} , the volume ratio

$$r := \frac{\text{Vol}(\mathcal{G} \cap \mathcal{C})}{\text{Vol}(\mathcal{C})},$$

can be regarded as a measure of the probability of correctness of the conjecture. A classical method to estimate r is the Monte Carlo approach, based on the generation of N independent identically distributed (i.i.d.) random samples within \mathcal{C} , which we denote by Δ^i , $i = 1, \dots, N$. An estimate of r based on the empirical outcomes of the N instances of the problem is given by $\hat{r}(N) = \frac{1}{N} \sum_{i=1}^N I_{\mathcal{G} \cap \mathcal{C}}(\Delta^i)$ where $I_{\mathcal{G} \cap \mathcal{C}}(\Delta^i) = 1$ if $\Delta^i \in \mathcal{G} \cap \mathcal{C}$ and 0 otherwise.

By the laws of large numbers for empirical probabilities, we can expect that $\hat{r}(N) \rightarrow r$ as $N \rightarrow \infty$. Probability inequalities for finite sample populations, such as the classical Chernoff bound [24], provide a lower bound N such that the empirical mean $\hat{r}(N)$ differs from the true probability r less than ϵ with probability greater than $1 - \delta$, i.e. $\text{Pr}\{|r - \hat{r}(N)| < \epsilon\} > 1 - \delta$, for $0 < \epsilon, \delta < 1$. The Chernoff bound is given by

$$N > \frac{1}{2\epsilon^2} \log\left(\frac{2}{\delta}\right). \quad (6)$$

Notice that the sample size N , given by (6), is independent on the size of \mathcal{B} and on the distribution.

To obtain an empirical estimate of r through execution of numerical experiments in our specific problem, the predicate can be modified in the finitely computable form

$$\mathbf{P}'_{GR}(G_0, G_f) = \{J(G_0, G_f) \leq \gamma\},$$

where $J(G_0, G_f)$ denotes the time employed by the last agent to reach its goal, and γ is a threshold to be suitably fixed.

An exhaustive probabilistic verification of the conjecture for wide ranges of all the involved variables remains untractable. To provide a meaningful set of results, however, some of the experimental parameters can be fixed according to criteria indicating the complexity of problems. In other terms, for a given size of the workspace \mathcal{B} , the safety distance d_s and the number of agents n can be chosen so that

- 1) the area occupied by the agents and their reserved discs is a significant portion of the available workspace, and
- 2) the average worst arrival time of agents is substantially larger than the time necessary for a solution computed disregarding collision avoidance.

The second criterion provides a qualitative information on the amount of deviations from nominal paths caused by collisions, hence on the amount of conflicts occurred.

Several experiments have been conducted to assess how these two indicators vary with the parameters (see Fig. 8). With the choice $\mathcal{B} = ([0, 800] \times [0, 700] \times [0, 2\pi])^{2n}$, $d_s = 18$ and $n = 10$, the area occupied by agents is 7% of the workspace, and the average worst arrival time is 80% longer than the unconstrained solution time. Another set of preliminary experiments have been conducted to choose a threshold time γ which was computationally manageable, yet sufficiently long not to discard solutions. The percentage of successes of the policy as a function of the threshold γ is reported in figure 8.

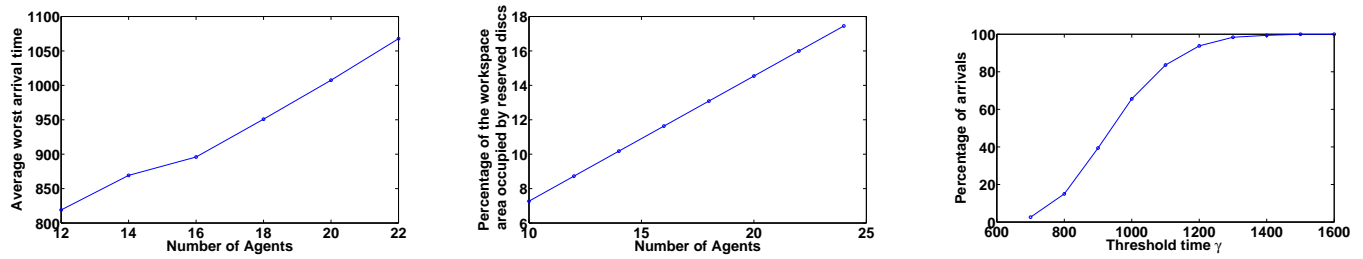


Fig. 8. Left: Average worst arrival time (over 300 experiments) vs. safety distance, for a system of 10 agents. The average unconstrained solution time is close to 520. Center: Percentage of workspace area occupied by agents and their reserved discs for different numbers of agents. Right: Percentage of arrivals with respect to threshold time γ .

From results obtained, it appears that only minor modifications of the outcomes should be expected for thresholds above $\gamma = 1600$. Finally, an estimate of the ratio r has been obtained by the probabilistic approach previously described. In order to have accuracy $\epsilon = 0.01$ with 99% confidence ($\delta = 0.01$), it was necessary by (6) to run 27000 experiments, with initial and final conditions uniformly distributed in the configuration space \mathcal{C} . Samples were generated by a rejection method applied to uniform samples generated in \mathcal{B} . None of these 27000 experiments failed to find a solution within time $\gamma = 4000$, hence $\hat{r}(N) = 1$. Hence, we can affirm with 99% confidence that the sparsity condition is sufficient to guarantee admissible plans for the generalized roundabout policy to within an approximation of 1% in case of $n = 10$ agents with safety disc of diameter $d_s = 18$.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we have considered a decentralized cooperative control policy for conflict resolution for multiple non-holonomic vehicles. Conditions on admissibility of problems for the policy to provide correct solutions have been investigated. A probabilistic method has been used to verify the correctness of a conjectured condition.

VI. ACKNOWLEDGMENT

Work done with partial support by EC contract IST-2004-004536 (IP RUNES), EC contract IST-2004-511368 (NoE HYCON), and NSF Career grant 0513058. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author and do not necessarily reflect the views of the supporting organizations.

REFERENCES

- [1] L. E. Dubins, "On curves of minimal length with a constraint on average curvature and with prescribed initial and terminal positions and tangents," *American Journal of Mathematics*, vol. 79, pp. 497–516, 1957.
- [2] S. M. LaValle and S. A. Hutchinson, "Optimal motion planning for multiple robots having independent goals," *IEEE Trans. on Robotics and Automation*, vol. 14, no. 6, pp. 912–925, 1998.
- [3] J. Peng and S. Akella, "Coordinating multiple robots with kinodynamic constraints along specified paths," in *Proc. Fifth International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, Nice, France, 2002.
- [4] B. P. Gerkey and M. J. Mataric, "Sold!: Auction methods for multi-robot coordination," *IEEE Trans. Robotics and Automation*, vol. 18, no. 5, pp. 758–768, 2002.
- [5] H. Chitsaz, J. O'Kane, and S. LaValle, "Exact pareto-optimal coordination of two translating polygonal robots on an acyclic roadmap," in *Proceedings of the International Conference on Robotics and Automation*, vol. 4, 2004, pp. 3981 – 3986.
- [6] V. J. Lumelsky and K. R. Harinarayan, "Decentralized motion planning for multiple mobile robots: the cocktail party model," *Autonomous Robots*, vol. 4, no. 1, pp. 121–35, 1997.
- [7] E. Klavins, "Communication complexity of multi-robot systems," in *Proc. Fifth International Workshop on the Algorithmic Foundations of Robotics*, Nice, France, 2002.
- [8] G. Inalhan, D. M. Stipanovic, and C. J. Tomlin, "Decentralized optimization, with application to multiple aircraft coordination," in *Proc. IEEE Conf. on Decision and Control*, 2002.
- [9] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Trans. Aut. Control*, vol. 48, no. 6, pp. 988–1001, 2003.
- [10] H. Tanner, A. Jadbabaie, and G. Pappas, "Flocking agents with varying interconnection topology," 2004, submitted to *Automatica*.
- [11] R. Olfati-Saber, "Flocking for multi-agent dynamic systems: Algorithms and theory," 2004, submitted to the *IEEE Trans. on Automatic Control* (Technical Report CIT-CDS 2004-005).
- [12] S. Loizu, D. Dimarogonas, and K. Kyriakopoulos, "Decentralized feedback stabilization of multiple nonholonomic agents," in *Proceedings of the International Conference on Robotics and Automation*, vol. 3, 2004, pp. 3012 – 3017.
- [13] C. Tomlin, G. J. Pappas, and S. Sastry, "Conflict resolution for air traffic management: A case study in multi-agent hybrid systems," *IEEE Trans. Aut. Control*, vol. 43, pp. 509–521, 1998.
- [14] R. Ghosh and C. J. Tomlin, "Maneuver design for multiple aircraft conflict resolution," in *Proc. American Control Conf.*, Chicago, IL, 2000.
- [15] C. Tomlin, I. Mitchell, and R. Ghosh, "Safety verification of conflict resolution maneuvers," *IEEE Trans. Intelligent Transportation Systems*, vol. 2, no. 2, 2001.
- [16] L. Pallottino, V. Scordio, and A. Bicchi, "Decentralized cooperative conflict resolution among multiple autonomous mobile agents," in *Proceedings of the Conference on Decision and Control*, vol. 5, Dec. 2004, pp. 4758–4763.
- [17] E. Frazzoli, L. Pallottino, V. G. Scordio, and A. Bicchi, "Decentralized cooperative conflict resolution for multiple nonholonomic vehicles," in *Proc. AIAA Guidance, Navigation and Control conference*, 2005.
- [18] R. Tempo, G. Calafiore, and F. Dabbene, *Randomized Algorithms for Analysis and Control of Uncertain Systems*, ser. Communications and Control Engineering. Springer-Verlag, 2003.
- [19] R. Alur, C. Courcoubetis, N. Halbwachs, T. Henzinger, P. Ho, X. Nicollin, A. Olivero, J. Sifakis, and S. Yovine, "The algorithmic analysis of hybrid systems," *Theoretical Computer Science*, vol. 138, pp. 3–34, 1995.
- [20] M. S. Branicky, "Studies in hybrid systems: Modeling, analysis, and control," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, 1995.
- [21] J. Lygeros, "Hierarchical hybrid control of large scale systems," Ph.D. dissertation, University of California, Berkeley, CA, 1996.
- [22] R. Alur, T. Dang, J. Esposito, Y. Hur, F. Ivančić, V. Kumar, I. Lee, P. Mishra, G. J. Pappas, and O. Sokolsky, "Hierarchical modeling and analysis of embedded systems," *Proceedings of the IEEE*, vol. 91, no. 1, pp. 11–28, 2003.
- [23] N. Lynch, R. Segala, and F. Vandraager, "Hybrid I/O automata revisited," in *Hybrid Systems IV: Computation and Control*, ser. Lecture Notes in Computer Science, M. Di Benedetto and A. Sangiovanni-Vincentelli, Eds. Springer-Verlag, 2001, vol. 2034, pp. 403–417.
- [24] H. Chernoff, "A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations," *Annals of Mathematical Statistics*, vol. 23, pp. 493–507, 1952.