

Decentralized cooperative policy for conflict resolution in multi-vehicle systems

Lucia Pallottino, Vincenzo G. Scordio, Emilio Frazzoli and Antonio Bicchi

Index Terms—Distributed control, Mobile robots, Mobile robot motion-planning

Abstract—In this paper we propose a novel policy for steering multiple vehicles between assigned independent start and goal configurations and ensuring collision avoidance. The policy rests on the assumption that agents are all cooperating by implementing the same traffic rules. However, the policy is completely decentralized, as each agent decides its own motion by applying those rules only on locally available information, and totally scalable, in the sense that the amount of information processed by each agent and the computational complexity of the algorithms are not increasing with the number of agents in the scenario. The proposed policy applies to systems in which new vehicle may enter the scene and start interacting with existing ones at any time, while others may leave. Under mild conditions on the initial configurations, the policy is shown to be safe, i.e. to guarantee collision avoidance throughout the system evolution. In the paper, conditions are discussed on the desired configurations of agents under which the ultimate convergence of all vehicles to their goals can also be guaranteed. To show that such conditions are actually necessary and sufficient, which turns out to be a challenging liveness verification problem for a complex hybrid automaton, we employ a probabilistic verification method. The paper finally reports on simulations for systems of several tens of vehicles, and with some experimental implementation showing the practicality of the approach.

I. INTRODUCTION

One of the challenging new applications foreseen for distributed control over networks is traffic management of multiple vehicle systems. The problem is posed in several different domains (such as e.g. traffic control of autonomous or semi-autonomous vehicles, on ground, air, or sea, planetary exploration, surveillance etc.) and with many different flavours (e.g. collision avoidance for cars at intersections, free-cruise assistance for small boats near harbors or aircraft near airports, formation flight for autonomous unmanned aerial or underwater vehicles, etc.). Multi agent systems offer many potential advantages with respect to single-agent systems such as speedup in task execution, robustness with respect to failure of one or more agents, and scalability. On the other hand, they introduce challenging issues such as the handling of distributed information data, the coordination among agents, the choice of communication protocols, the design and verification of decentralized control laws, and security issues.

Lucia Pallottino, Vincenzo G. Scordio and Antonio Bicchi are with the Dept. Electrical Systems and Automation, and the “Enrico Piaggio” Interdept. Research Center of the University of Pisa, via Diotisalvi, 2, 56100 Pisa, Italy

Emilio Frazzoli is with the Mechanical and Aerospace Engineering Department, University of California at Los Angeles, 37-144 Engineering IV, 420 Westwood Plaza Los Angeles, CA, 90095

Manuscript submitted July, 2006. Corresponding author: Lucia Pallottino, l.pallottino@ing.unipi.it

This paper focuses on the problem of collision avoidance management of a very large and dynamically changing number of autonomous vehicles moving in a shared environment. Each vehicle is assumed to have different task to accomplish in terms of reaching a target configuration. In a centralized approach a single decision maker must know current and desired configuration of all agents in order to determine collision free controls for each vehicle. Although correct and complete algorithms for the centralized traffic management problem may exist, they typically require a large amount of computational resources. Furthermore, centralized approaches typically are very prone to faults of the decision maker. Decentralized approaches require that each vehicle plans its own trajectory based only on information limited to neighboring vehicles. A decentralized approach is typically faster to react to unexpected situations, but safety verification is an issue as domino effects of possible conflicts may prevent convergence to solutions in some conditions. Furthermore, a decentralized approach ensures scalability of the system. Indeed, the amount of information required by each vehicle is independent of the total number of vehicles in the scenario that may change in time. In particular, in our approach vehicles are aware of the position and orientation of nearby agents, within a certain sensing or communication radius, but have access to no other information such as goals or 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.

Traffic management has been often attacked in the hypothesis that vehicles have rather simple dynamics, allowing them to stop rapidly to clear possible impending conflicts, and change direction of motion instantly. This assumption is however inapplicable (or imposes too conservative limits on velocities) for most practical vehicles, which have important dynamics preventing immediate stops (as e.g. with cars or marine vessels) or not allowing stops at all (as with aircraft). In this paper, we focus on the harder problem involving agents which cannot stop, by considering a simplified but realistic kinematic models of their dynamics. The model assumes that the vehicle move with constant speed subject to curvature bounds. This model well suites the scenario of vehicles such as aircrafts cruising on a planar airspace.

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., [1]–[4]; a characterization of Pareto-optimal solutions has been provided in [5]. In [6] the problem of path planning is divided into global and local path planning, and AI techniques are used in combination with real-time techniques. In [7] and [8], formations of robots are considered, where a motion plan for the overall formation is used to control a single “lead” robot while the “followers” are governed by local control laws, sensing their positions relative to neighboring robots. In [9] a framework exploiting the advantages of centralized and decentralized planning for multiple mobile robots with limited ranges of sensing and communication maneuvering in dynamic environments, is presented. Several decentralized algorithms have appeared, e.g. [10], [11] for holonomic robots, and [12] for aircraft-like vehicles. The literature on flocking and formation flight, which has flourished recently (e.g., [13]–[15]), while ultimately leading to conflict-free collective motion, does not address individual objectives, and vehicles 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 [16]. However, the control laws in [16] are not directly applicable to our case, in which vehicles are constrained to move at constant speed, and cannot stop or back up. Furthermore, they assume that each agent is aware of the (fixed) total number of agents in the scenario.

In the literature dealing more specifically with air traffic control, the early work of [17] 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 [18], [19]. 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 [20]. 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 [21] which is (i) spatially decentralized, and (ii) provably safe, regardless of the number of vehicles present in the environment. The method builds on [10], 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.

Our policy will be proved to be safe for an arbitrarily large number of vehicles, and indeed very effective in negotiating conflicts of several tens of vehicles. 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. Notice that in [10], no liveness guarantees were given, and indeed counterexamples were provided.

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 [22]. 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

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

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

Define the map $d : SE(2) \times SE(2) \rightarrow \mathbb{R}^+$ as the distance between the positions of two agents; in coordinates,

$$d(g_1, g_2) = \|(x_1, y_1) - (x_2, y_2)\|_2.$$

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 , i.e., if $d(g_i(t_c), g_j(t_c)) < d_s$. 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 is a map $\pi : Z \times 2^{SE(2)} \rightarrow [-1, 1]$, $(z, \bar{g}) \mapsto \omega$ that associates to an individual agent a control input, based on a set of locally-available *internal variables* $z \in Z$, and on the current configuration of other agents in the environment. We use the shorthand $\bar{g} \subset \{g_1, \dots, g_n\} \subset SE(2)$, to indicate a set of cardinality $\text{card}(\bar{g}) < n$, summarizing the available information about other agents. 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; that is, we say that policy π is spatially decentralized if

$$\pi(z, \bar{g}) = \pi(z, \text{Neigh}(g, \bar{g}, d_a)),$$

where the map Neigh extracts neighbors of g from \bar{g} , i.e., $\text{Neigh}(g, \bar{g}, d_a) = \{\hat{g} \in \bar{g} | d(g, \hat{g}) \leq d_a\}$.

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.

The objective of this paper is to report a spatially decentralized feedback control policy that satisfies, under certain conditions, the following two properties:

- **Safety:** No conflicts are generated, i.e.,

$$\forall t > 0, i, j \in \{1, \dots, n\}, i \neq j : d(g_i(t), g_j(t)) \geq d_s. \quad (2)$$

- **Liveness:** At least one vehicle eventually reaches its destination:

$$\exists t_f \geq 0, i \in \{1, \dots, n\} : g_i(t_f) = g_{f,i}. \quad (3)$$

Note that if agents are removed from the environment upon arrival to their target, the liveness condition stated above can be applied recursively, to ensure that all agents will eventually reach their targets.

A basic assumption is that decisions are taken by each agent according to an a priori set of rules that each vehicle respects, i.e. a cooperative behavior between agents is required.

III. THE PROPOSED MOTION COORDINATION POLICY

In this section, we propose a spatially decentralized feedback control policy based on a number of discrete *modes of operation*, and as such the closed-loop system is a hybrid system. Properties of the control policy will be described in Section IV. In order to introduce our control policy, we need to define some of its elements.

A. Reserved region

This policy is based on a concept of *reserved region*, over which each active agent claims exclusive ownership. Let the map $c : SE(2) \rightarrow \mathbb{R}^2$, $(x, y, \theta) \mapsto (x^c, y^c)$ associate to the configuration of an agent the center of the circle it would describe under the action of a constant control input $\omega = -1$. In other words,

$$(x^c, y^c) = c(x, y, \theta) = (x + \sin(\theta), y - \cos(\theta));$$

see figure 1.

The reserved region for the i -th agent is defined as a disc of radius $1 + R_s$ centered at $c(g_i)$:

$$R_i(t) = \{(x, y) \in \mathbb{R}^2 : \|(x, y) - c(g_i(t))\|_2 \leq 1 + R_s\}. \quad (4)$$

The motion of the point (x_i^c, y_i^c) is described by the following equations:

$$\begin{aligned} \dot{x}_i^c(t) &= (1 + \omega_i(t)) \cos \theta_i(t) \\ \dot{y}_i^c(t) &= (1 + \omega_i(t)) \sin \theta_i(t). \end{aligned} \quad (5)$$

Furthermore, we associate a heading angle to the reserved disc that coincides with the agent heading θ_i . Hence, $\dot{\theta}_i^c(t) = \omega_i(t)$. Our policy is based on the following basic observations: the model described by (1) and (5) is such that the reserved region (i) can be stopped at any time, by setting $\omega = -1$,

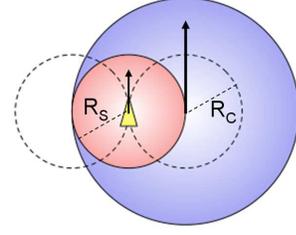


Fig. 1. The reserved region of a nonholonomic vehicle.

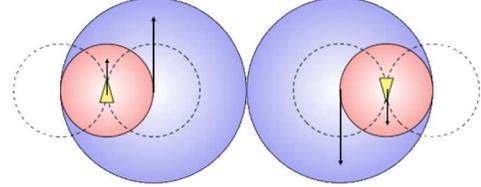


Fig. 2. Worst case scenario for the choice of the alert distance. In our case $R_C = 1$ and $R_S = \frac{d_s}{2}$.

and (ii) once stopped, it can be moved in any direction, provided one waits long enough for the heading θ to reach the appropriate value. As a consequence, for example, the center of the reserved region can follow any continuous path within an arbitrarily small tolerance, unlike model (1). Note that it is always possible to keep the reserved region fixed, with the corresponding agents moving along a minimum-radius circle entirely contained within it, see Figure 1.

B. Constraints

A sufficient condition to ensure safety is that the interiors of reserved regions are disjoint at all times; if such a condition is met, conflicts can be avoided if agents hold their reserved regions fixed, and move within them (by setting $\omega = -1$). As a consequence, each point of contact between reserved regions defines a constraint on further motion for both agents involved. More precisely, if the reserved region of agent i is in contact with the reserved regions of agents with indices in $J_i \subset \{1, \dots, n\}$, the motion of the agents is constrained as follows

$$\dot{x}_i^c(x_i^c - x_j^c) + \dot{y}_i^c(y_i^c - y_j^c) \geq 0, \quad \forall j \in J_i. \quad (6)$$

In other words, the velocity of the i -th reserved region is constrained to remain in the convex cone determined by the intersection of a number of closed half-planes.

Note that the full set of constraints can be computed assuming that each agent is aware of the configuration of all agents within an alert distance $d_a = 2(1 + 2R_s)$; see to figure 2. In addition, the amount of information needed to compute the bound is uniformly bounded, independent from the total number of agents in the system: in fact, the maximum number of agents whose reserved region is in contact with the reserved region of the computing agents is six, see figure 3. Concluding the proposed policy does not depend upon the number of the agents in the environment.

Let us define the set-valued map $\Theta : SE(2) \times 2^{SE(2)} \rightarrow 2^{S^1}$, associating to the configurations of an agent g and of its

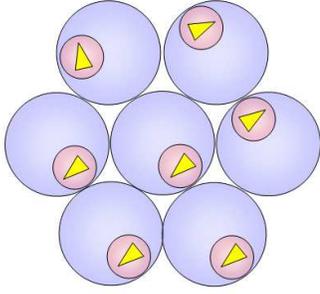


Fig. 3. The vehicle in the center of the cluster exchange information with the six neighbor agents.

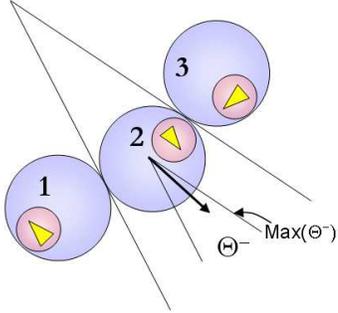


Fig. 4. The set of allowable directions in which the center of the i -th reserved region can move, generated by the contact with the reserved regions of vehicles j , m and k respectively.

neighbors (in \bar{g}) the set of allowable directions in which the reserved region of the computing agent can translate without violating the constraints (6). For a connected, non-empty set $B \subset S^1$, $B \neq \emptyset$, let us define $\max(B)$ and $\min(B)$ as the elements on the boundary of B , respectively in the positive and negative direction with respect to the bisectrix of B . Finally, define the map $\Theta^-(g, \bar{g}) = \Theta(g, \bar{g}) \setminus \min(\Theta(g, \bar{g}))$. In other words, the output of Θ^- is an open set, obtained removing the boundary in the clockwise direction of the cone of feasible reserved region translations, see figure 4. Whenever Θ is a proper subset of S^1 , $\max(\Theta)$, $\min(\Theta)$, and Θ^- are well defined. If $\Theta = \emptyset$, or $\Theta = S^1$, we set $\Theta^- = \Theta$.

C. Holding

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

D. Right-turn-only steering policy

Our concept for decentralized conflict-free coordination is based on maintaining the interiors of reserved regions disjoint. Assuming that no constraints are violated, an agent will attempt to steer the center of its own reserved region towards the position it would assume at the target configuration. In a free environment, this can be accomplished switching between the hold state and a straight state:

$$\omega = \begin{cases} 0 & \text{if } \|\Delta_f\|_2 > 0 \text{ and } \theta = \phi(\Delta_f) \\ -1 & \text{otherwise} \end{cases} \quad (7)$$

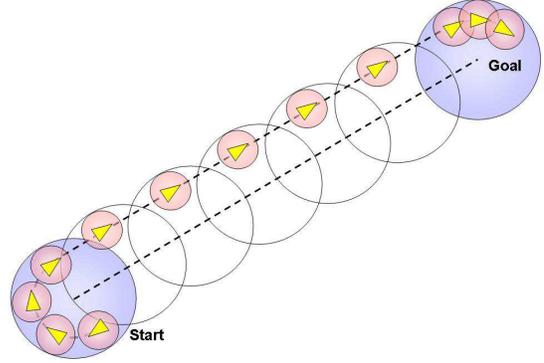


Fig. 5. The proposed policy in a single agent case.

where $\Delta_f = c(g_f) - c(g)$, and $\phi : \mathbb{R}^2 \setminus 0 \rightarrow S^1$ is a function returning the polar angle of a vector. Note that reserved region move along straight lines according to (7); clearly, such a policy is not optimal (in a minimum-time or minimum-length sense), but it does provide a simple feasible path for the agent from the current configuration to its target, see figure 5.

E. Rolling on a stationary neighboring reserved region

If the path of the reserved region to its position at the target is blocked by another reserved region, 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 region. Since in our setup agents communicate only information on their states, not on their future intentions, care must be exercised in such a way that the interiors of reserved regions remain disjoint

Let us start by assuming that the reserved region of the neighboring agent remains stationary; in order to roll on such region, without violating safety constraints, the control input must be set to

$$\omega = \begin{cases} (1 + d_s/2)^{-1} & \text{if } \Theta^-(g, \bar{g}) \neq \emptyset \text{ and } \theta = \max(\Theta) \\ -1 & \text{otherwise} \end{cases} \quad (8)$$

A possible trajectory in case of a stationary obstacle is reported in figure 6.

The above policy is obtained by switching between the hold state and a roll state; note that when in the roll state, the agent is not turning at the maximum rate.

Note that (8) also addresses the case in which the agent's motion is constrained by more than one contact with other agents' reserved regions. The only case in which the agent will not transition to the roll state, is the degenerate case in which Θ is a singleton, and Θ^- is empty.

F. Non-stationary neighbors

In general, the reserved region 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 regions of two or more agents executing (8) will always remain disjoint, it is possible that contact between two agents is lost unexpectedly (recall that the control input of other agents, their

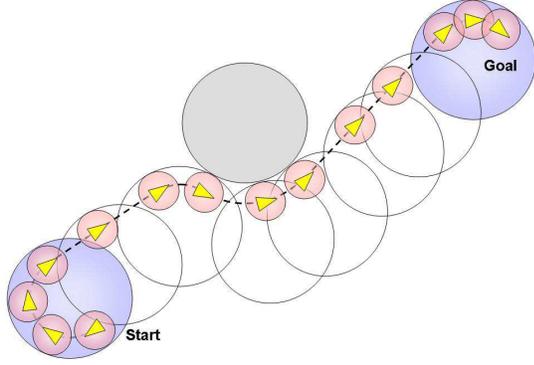


Fig. 6. Trajectory in case of stationary neighboring reserved region

constraints, and their targets, are not available). 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`.

G. 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 now introduce the hybrid system modeling the dynamics of a single agent. We define a hybrid system as a tuple

$$\mathcal{S} = (Q, X, U, \Phi, \Delta, \text{Inv}, \text{Init}),$$

where Q is a set of discrete states, X is the continuous state space, U is a set of exogenous inputs, $\Phi : Q \times X \times U \rightarrow TX$ is a function describing the continuous dynamics of the system, Δ is a relation describing discrete transitions, and Inv, Init denote the invariant and initial conditions set, respectively. We refer the reader to the relevant literature for a more in-depth discussion of the hybrid systems formalism (e.g., [23]–[26] and references therein).

More in detail, the model for an individual agent can be specified as follows:

$$\mathcal{A} = (\{\text{roll}, \text{roll2}, \text{hold}, \text{straight}\}, SE(2) \times \mathbb{R}, 2^{SE(2)}, \Phi_{\text{GR}}, \Delta_{\text{GR}}, \text{Inv}_{\text{GR}}, \text{Init}),$$

- The discrete states $Q = \{\text{roll}, \text{roll2}, \text{hold}, \text{straight}\}$ correspond to constant inputs $\omega_{\text{roll}} = (1 + d_s/2)^{-1}$, $\omega_{\text{roll2}} = +1$, $\omega_{\text{hold}} = -1$, and $\omega_{\text{straight}} = 0$, respectively.
- $X = SE(2) \times \mathbb{R}$: in addition to its own configuration, each agent can keep track of time through a clock τ .

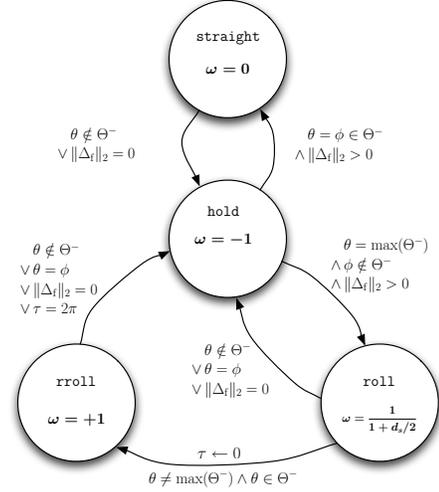


Fig. 7. A hybrid automaton describing the Generalized Roundabout policy.

- $U = 2^{SE(2)}$: The exogenous input is a set $\bar{g} \subset SE(2)$ summarizing the available information about other agents. Since we are dealing with a decentralized policy, \bar{g} can be restricted to contain solely neighbors within an alert distance $d_a = 4 + d_s$.
- The map Φ_{GR} is derived from (1), substituting the appropriate value for ω , based on the discrete mode, and by the clock rate $\dot{\tau} = 1$, i.e., it can be written in coordinates as follows:

$$\begin{aligned} \dot{x} &= \cos(\theta) \\ \dot{y} &= \sin(\theta) \\ \dot{\theta} &= \omega_q, \quad q \in Q \\ \dot{\tau} &= 1. \end{aligned} \quad (9)$$

- The initial set for each agent is unrestricted, i.e., the hybrid state $z = (q, x)$ can take any value in $Q \times X$.
- We do not explicitly write down the GR policy and its transition relations, guards, and invariants, but we refer the reader to Figure 7, which should provide the necessary detail in a clearer fashion.

The multiple-vehicle system we are considering is the parallel composition of n agents

$$\mathcal{S}_{\text{GR}} = \mathcal{A}_1 | \mathcal{A}_2 | \dots | \mathcal{A}_n \quad (10)$$

coupled through their configurations, communicated through the individual agents' exogenous inputs; \mathcal{S}_{GR} does not have exogenous inputs itself (We do not define the operation of parallel composition here; see, e.g., [27] for details.)

IV. 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. 8. In this section, we investigate properties of the proposed policy and 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*).

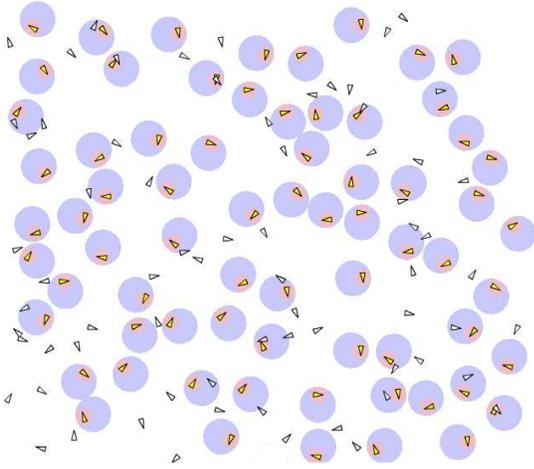


Fig. 8. A conflict resolution problem with 70 agents in narrow space, for which the proposed policy provides a correct solution. Initial configurations are identified by the presence of gray circles, indicating their reserved discs.

A. Admissibility

Consider a framework in which new agents may issue a request to enter the scenario at an arbitrary time and with an arbitrary “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 $\neg \mathbf{P}_1(G(t)) \wedge \neg \mathbf{P}_2(G_f)$. Simple tests to check the two properties are needed for the Generalized Roundabout Policy.

B. Well-posedness

The first step in our analysis of S_{GR} is to verify that the Generalized Roundabout Policy leads to a well posed dynamical system, i.e., a solution exists and is unique, for all initial conditions within a given set. Indeed,

Theorem 1: The hybrid system S_{GR} is well posed, for all initial conditions in which the interiors of reserved disks are

disjoint, i.e., $\|c(g_i) - c(g_j)\| \geq 2 + d_s, \forall i, j \in \{1, \dots, n\}$.

Proof: The map (9) is globally Lipschitz in the state and in the control input; moreover, control inputs are constant within a discrete mode. The parallel composition of n copies of the continuous dynamics (9) is also globally Lipschitz. Hence, in order to establish well posedness of it is sufficient to show that there is no accumulation point of switching times, i.e., the number of switches in an open time interval is bounded, and the control input signal ω is piecewise continuous.

First of all, note that the number of instantaneous switches is bounded by three: the specification of invariant conditions in Figure 7 prevents infinite loops without time advancement. This can be verified by inspection of the invariants.

Let t_0 denote the time at which a switch in the discrete state has occurred, we need to show that there exists a $t' > t_0$ such that there are no switches in the open interval (t_0, t') . For simplicity, assume that the discrete state at time t_0 is the terminal state of the sequence of instantaneous switches occurring at t_0 .

In the following, we will consider the i -th agent, and compute bounds on the time separation between switches, based on the current state of all agents. We have the following cases:

a) *Case 1: $q_i(t_0) = \text{hold}$* : A switch can be triggered by the following:

- $\Delta_{f,i} = 0, \theta_i = \theta_{f,i}$: The agent reaches its final configuration, and is removed from the system.
- $\Delta_{f,i=0} > 0, \theta_i = \phi_i \in \Theta_i^-$: the agent transitions to the discrete state `straight`.
- $\Delta_{f,i} > 0, \theta_i = \max(\Theta_i)$: the agent transitions to the discrete state `roll`.

None of the three above events can occur in the time interval $(t_0, t_0 + \delta_{1,i})$, with $\delta_{1,i} = \min\{\theta_{f,i} - \theta_i, \phi_i - \theta_i, \max(\Theta_i) - \theta_i\}$; the angle differences are meant to be counted in the direction of angular motion of the agent, modulo 2π .

b) *Case 2: $q_i(t_0) = \text{straight}$* : A switch can be triggered by the following:

- $\Delta_{f,i} = 0$: the reserved disk has been steered to its final configuration, and the agent transitions to the discrete state `hold`.
- $\phi_i \notin \Theta_i^-$: a new constraint on the motion of the reserved disk is activated, as the consequence of a contact with another agent’s reserved disk

Neither of the two above events can occur in the time interval $(t_0, t_0 + \delta_{2,i})$, with $\delta_{2,i} = \min\{\|\Delta_{f,i}\|_2, \min_{j \neq i} \{\|c(g_i) - c(g_j)\|_2 - (2 - d_s)\}\}$.

c) *Case 3: $q_i(t_0) = \text{roll2}$* : A switch can be triggered by events that have already been considered above, plus the time-out condition $\tau < 2\pi$. Hence no switches can occur in time interval $(t_0, t_0 + \delta_{3,i})$, where $\delta_{3,i} = \min\{\delta_1, \delta_2, 2\pi - \tau\}$.

d) *Case 4: $q_i(t_0) = \text{roll}$* : This is the only delicate case, as instantaneous transitions can be triggered by other agents’ actions. Let us indicate with j the index of the agent generating the constraint corresponding to $\max(\Theta_i)$. If $q_j = \text{hold}$, then the invariant $\theta_i = \max(\Theta_i)$ is preserved as the reserved disk of the i -th agent rolls on the reserved disk of

the j -th agent; switches can be triggered by events considered above. If $q_j \neq \text{hold}$, the reserved disks of the two agents will detach at time zero—thus triggering a transition of the discrete state of the i -th agent to `roll2`; however, since the motion of the j -th agent is constrained by agent i , in such a way that the envelope of the reserved disk of agent j forms an angle $\chi_{ij} > 0$ (since Θ_j^- has been defined as an open set), the time at which the next switch can occur in this case is no sooner than $t_0 + 2 \sin(\chi_{ij}/2)$. Hence, an additional switch cannot happen in the interval $(t_0, \delta_{4,i})$, with $\delta_{4,i} = \min\{\delta_2, (\phi_i - \theta_i)(1 + d_s/2), 2 \sin(\chi_{ij}/2)\}$.

Summarizing, for the whole system, if t_0 is a switching time for at least one of the agents, no other agents can switch within the interval $(t_0, t_0 + \delta)$, where $\delta = \min_i\{\delta_{1,i}, \delta_{2,i}, \delta_{3,i}, \delta_{4,i}\} > 0$.

Therefore, in the following, only initial configurations with disjoint reserved discs are considered.

C. Safety

A test for property \mathbf{P}_1 is provided by the following theorem.

Theorem 2: If the reserved disks of at least two agents in G overlap, property $\mathbf{P}_1(G)$ is verified. In other words, if $\|c(g_i) - c(g_j)\| \geq 2 + d_s, \forall i, j \in \{1, \dots, n\}, i \neq j$, it follows that $\forall t \geq 0, d(g_i(t), g_j(t)) > d_s, \forall i \in \{1, \dots, n\}, j \neq i$.

Proof: The proof of the theorem follows directly from the fact that trajectories $g_i(t), i = 1, \dots, n$ are continuous functions of time. Moreover, within each state the feedback control policy has been chosen so that reserved discs never overlap: a transition is always enabled to the `hold` state, which stops the reserved disk instantaneously. Since the agents are always contained within their reserved disk, at a distance $d_s/2$ from its boundary, safety is ensured. ■

D. Liveness

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 $m = 2, \dots, n$,

$$\text{card}\{(x_{f,i}^c, y_{f,i}^c) \in P_f^c : \|(x_{f,i}^c, y_{f,i}^c) - (x, y)\|_2 < \rho(m)\} < m, \quad (11)$$

where

$$\rho(m) = \begin{cases} 2(1 + R_S) & \text{for } m \leq 4, \\ (1 + \cot(\frac{\pi}{m}))(1 + R_S) & \text{for } m \geq 4. \end{cases} \quad (12)$$

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.

Consider the property:

P₃: A target configuration set $G_f = \{g_{f,i}, i = 1, \dots, n\}$ is clustered if the sparsity condition (11) is violated.

Theorem 3 (Necessary conditions for liveness): Property $\mathbf{P}_2(G_f)$ is verified for the GR policy if $\mathbf{P}_3(G_f)$ is verified, i.e. $\mathbf{P}_3(G_f) \Rightarrow \mathbf{P}_2(G_f)$.

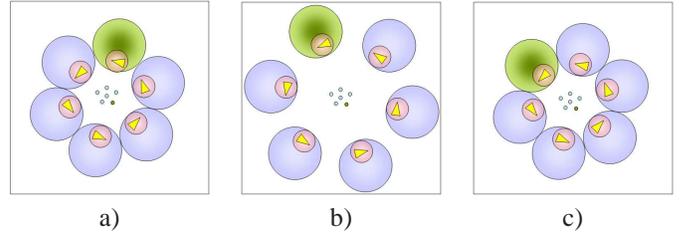


Fig. 9. Livelock-generating conditions for the GR policy with $\hat{m} = 6$.

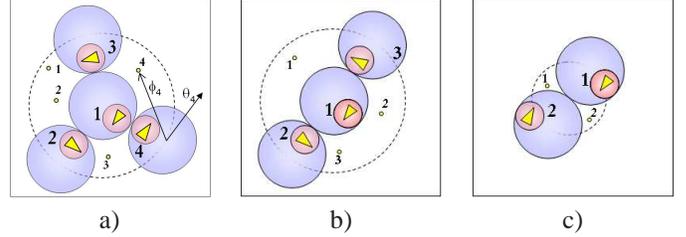


Fig. 10. Blocking executions of the GR policy with $\hat{m} \leq 4$.

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 (11), 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. 9-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 `roll2` (fig. 9-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. 9-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 tangent direction (see fig. 10-a and -b). Place the initial position of the reserved disc of the remaining agent in the center of $C_{\hat{m}}$. By applying the GR policy, this agent remains indefinitely in the `hold` state while the other $\hat{m} - 1$ 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 $\hat{m} = 2$. The construction and behaviour in this case is completely analogous to the case $\hat{m} \geq 5$ (see fig. 10-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 general proof of sufficiency

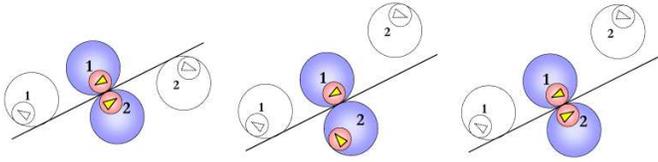


Fig. 11. Three possible situation for two agents with reserved discs in contact, agent 1 is such that $q_1 = \text{straight}$. On the left $q_2 = \text{straight}$, in the middle $q_2 = \text{hold}$, on the right $q_2 = \text{roll}$

appears to be very complex; however, we can now provide a demonstration of sufficiency in the simple case $n = 2$.

Theorem 4: Consider two vehicles such that the center of the reserved disc in final configurations are at distance larger than $2d_s + 4$. The GR policy allows the vehicles to reach their final destinations in finite time, from all initial conditions such that the interiors of the reserved disks are disjoint.

Proof: If the reserved disks of the two vehicles will never be tangent, the two vehicles will reach their goal with the sequence of controls $\omega = -1, \omega = 0, \omega = -1$. Otherwise, when a contact between the reserved discs occurs, we have six different combinations of mode of operation:

- Case 1 $q_1 = q_2 = \text{straight}$.
- Case 2 $q_1 = \text{straight}, q_2 = \text{hold}$.
- Case 3 $q_1 = \text{straight}, q_2 = \text{roll}$.

Those cases, reported in figure 11 are such that the contact will be immediately lost. In the first case no other contacts will be generated and the goals will be reached with a sequence of transitions `straight`, `hold` for both vehicles. In the second case the first vehicle will reach its final destination with a sequence `straight`, `hold`, while the second one will maintains control `hold` until it is no longer blocked by the first vehicle, and can move towards its goal; the reserved disks will no longer touch. In the third case, the second agent will transition to the `roll2` state as soon as contact is lost. The reserved disk of second agent will never be tangent again to the reserved of the first one. Hence, the second agent will reach its final destination with a sequence `roll2`, `hold`, `straight`, `hold`, or `roll2`,`straight`,`hold`, depending on the initial and final configurations.

Let us now consider the following cases reported in figure 12

- Case 4 $q_1 = q_2 = \text{hold}$.
- Case 5 $q_1 = \text{hold}, q_2 = \text{roll}$.

It is sufficient to discuss the second case, since if both vehicles are in state `hold` they will reach a configuration that is equivalent to the second case unless one of them can move through its final configuration without contacts of the reserved disks. If this occurs, one of the vehicles will be in state `straight` and this is the case 2 discussed above.

In the second case the second vehicle will turn on the left so that the second reserved discs will slide along the first one until one of the two vehicle are able to move through the goal or they reach the configuration of case 6 (that will be discussed below).

- Case 6 $q_1 = q_2 = \text{roll}$.

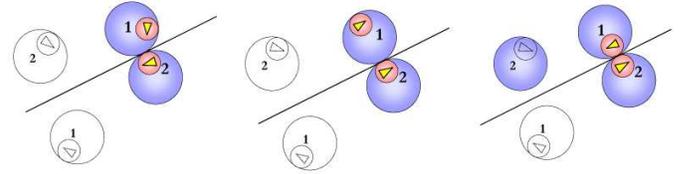


Fig. 12. Three possible situation for two agents with reserved discs in contact, on the left and in the middle, agent 1 is such that $q_1 = \text{hold}$, while agent 2 $q_2 = \text{hold}$ and $q_2 = \text{roll}$. On the right $q_1 = q_2 = \text{roll}$.

In this case the contact will be lost immediately, and both vehicles will switch to `roll2`; reserved disks may touch again. If a new contact occurs, the point of contact between the reserved discs has moved counterclockwise in the first vehicle's frame and clockwise on the second one. After this new contact both vehicles are in the `hold` state. If one of the vehicle can move through its final configuration by switching to the `straight` state, the configuration is equivalent to Case 2. Otherwise this procedure is repeated. But after enough time if the distance between target configurations is larger than $2d_s + 4$, one of the two vehicles will be able to move through its goal since at least one of the goals is not covered by the cluster movements. In this case for one vehicle $\omega = 0$ and the configuration is equivalent to one of the previous cases. ■

In the next section, we describe a method to approach the problem of sufficiency from a probabilistic point of view.

V. PROBABILISTIC VERIFICATION OF THE GR POLICY

Consider the following statement:

Conjecture [Sufficient conditions for admissibility]

The GR policy provides a non-blocking solution for all safe and non clustered plans (G_0, G_f) .

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.

The conjecture can be represented with the logic statement: $\neg \mathbf{P}_1(G(t)) \wedge \neg \mathbf{P}_3(G_f) \Rightarrow \mathbf{P}_{GR}(G_0, G_f) = \neg \mathbf{P}_1(G(t)) \wedge \neg \mathbf{P}_2(G_f)$

A probabilistic verification of the conjecture can be obtained following the approach described below (for more details, see e.g. [22]) as described in [28].

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 $\mathcal{G} = \{(G_0, G_f) \in \mathcal{B} | \mathbf{P}_{GR}(G_0, G_f)\}$ denote the ‘‘good’’ set of problem data for which the predicate applies. Also, let $\mathcal{C} = \{(G_0, G_f) \in \mathcal{B} | \neg \mathbf{P}_1(G_0) \wedge \neg \mathbf{P}_3(G_f)\}$ denote the set of safe and non clustered plans.

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 [29], 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. $Pr\{|r - \hat{r}(N)| < \epsilon\} > 1 - \delta$, for $0 < \epsilon, \delta < 1$. The Chernoff bound is given by

$$N > \frac{1}{2\epsilon^2} \log\left(\frac{2}{\delta}\right). \quad (13)$$

Notice that the sample size N , given by (13), 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. 13 and 14). 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 15. 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 (13) 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

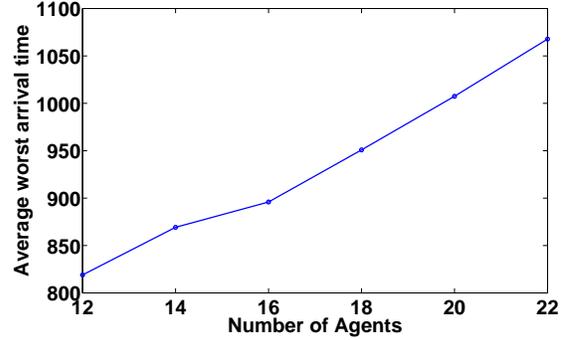


Fig. 13. Average worst arrival time (over 300 experiments) vs. safety distance, for a system of 10 agents. The average unconstrained solution time is close to 520.

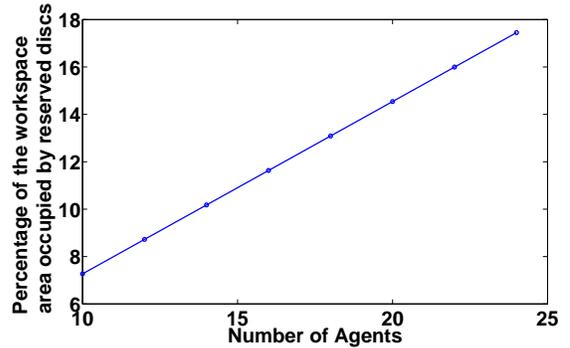


Fig. 14. Percentage of workspace area occupied by agents and their reserved discs for different numbers of agents.

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 $d_s = 18$.

A. Qualitative evaluation of the sparsity condition and of the liveness of the policy

We are now interested in providing qualitative evaluations of sparsity condition on the targets and the liveness of the chosen policy .

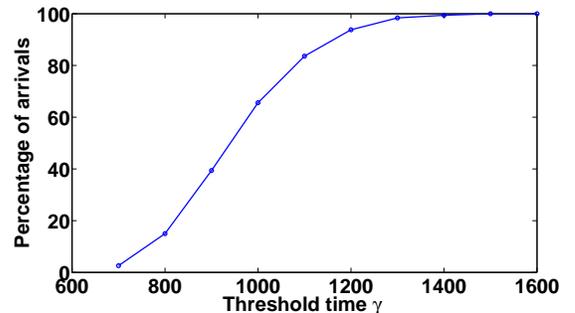


Fig. 15. Percentage of arrivals with respect to threshold time γ .

The dimension of \mathcal{C} in \mathcal{B} depends on the value of the number of agents n and the value of the associated safety radius R_S . Figure 16 represents the normalized dimension of \mathcal{C} in \mathcal{B}_n with respect to variation of $n \in \{2, \dots, 20\}$ and $R_S \in \{2, \dots, 40\}$. In figure 17 the z-axis view is reported. Projections of the isodimensional curves on the (n, R_S) plane appear to be hyperbolas, i.e. $n R_S = const..$

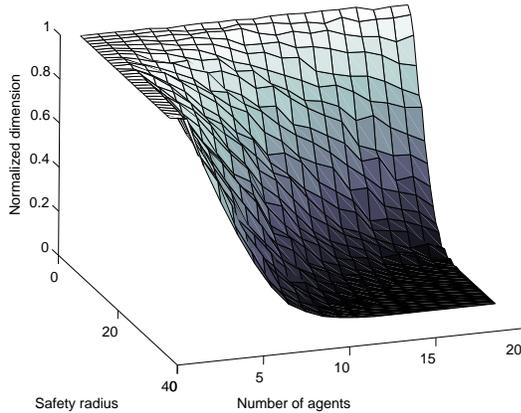


Fig. 16. The normalized dimension of \mathcal{C} in \mathcal{B} with respect to variation of n and R_S .

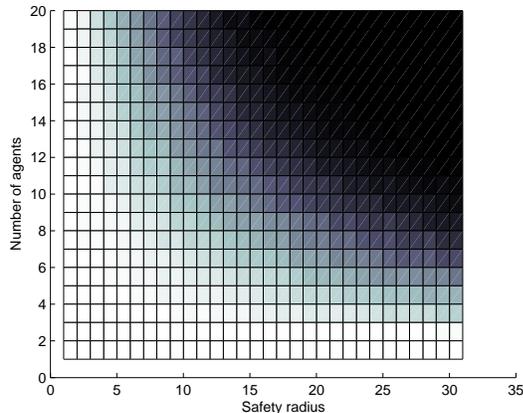


Fig. 17. Projections of the isodimensional curves on the (n, R_S) plane appear to be hyperbolas.

Using values of n and R_S such that the dimension of \mathcal{C} in \mathcal{B}_n is larger or equal to 95% we have verified, with the proposed probabilistic approach, that with 99% confidence the sparsity condition is sufficient to guarantee liveness of the generalized roundabout policy to within an approximation of 1%. For the remaining 5% of $\mathcal{B}_n \setminus \mathcal{C}$ more than 20000 simulations have been run. In the 96.433% of cases such simulations have terminated with the reaching of the goal configurations, i.e. no livelock has occurred. In conclusion, regarding the liveness property of the proposed Roundabout policy, we can affirm that for some particular values of n and R_S in more of $0.99 \cdot 0.95 + 0.96433 \cdot 0.05 = 99.8\%$ of cases all agents will eventually reach the goal configurations.

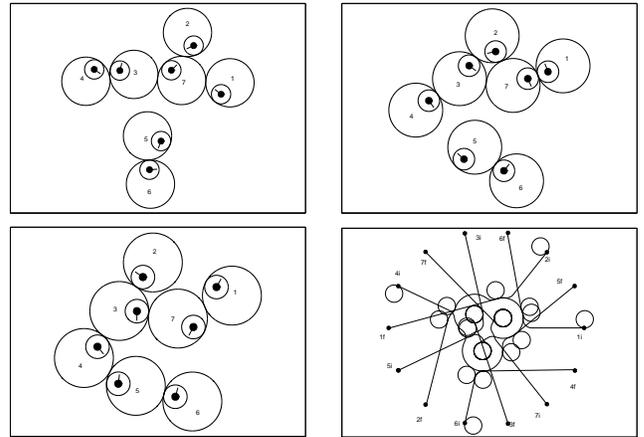


Fig. 18. Significant instants and whole agents trajectories (bottom-right) of a simulation with seven agents

Furthermore, notice that for those value of n and R_S , the total space occupied by agents is around the 4 – 5% of the whole workspace. To give an idea, in terms of agents occupancy this means that in a workspace of dimension 7meter \times 8meter we are able to manage safely 10 agents with a safety disc diameter of 60 centimeters.

VI. SIMULATIONS AND ARCHITECTURE IMPLEMENTATION

As reported in previous section, a large number of simulations have been conducted with parameters of different values. In figure 18 some significant instants and whole agents trajectories of a simulation with seven agents are reported for reader convenience.

On the other hand, a scalable platform has been designed for safe and secure decentralized traffic management of multi-agent mobile systems and applied to the proposed scenario with the GRP policy. The architecture is based on wireless communication between agents and provides vehicles of services such as the localization service or the authorization one. In the current implementation localization is provided to vehicles by a centralized server through camera that monitor the environment. The authorization server is responsible of testing the admissibility of a proposed plan as described in section IV-A. In figure 19 some screenshots of a three-vehicle case is reported with overprinted reserved discs.

VII. CONCLUSIONS AND FUTURE WORK

In this paper, we have outlined a novel spatially decentralized, cooperative policy for conflict-free motion coordination of non-holonomic vehicles. All of the computations involved in the proposed policy are spatially decentralized, and their complexity is bounded regardless of the number of agents, thus making the policy scalable to large-scale systems. The policy gives rise to a hybrid system, which can be shown to be well posed, and safe, if the initial conditions satisfy a rather non-restrictive (but possibly conservative) condition. Conditions on admissibility (safety and liveness) 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.

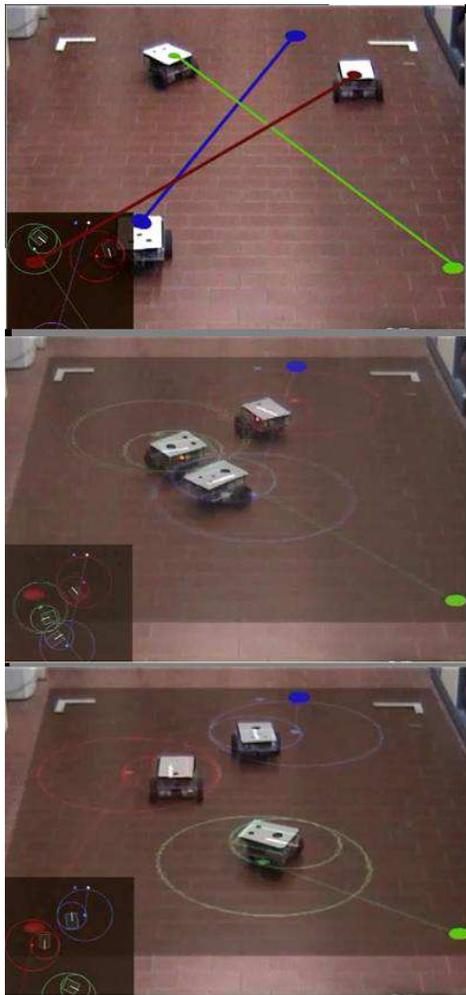


Fig. 19. Snapshots of GRP evolution in a three-vehicle scenario

VIII. 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] 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.
- [2] 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.
- [3] 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.
- [4] K. Azarm and G. Schmidt, "Conflict-free motion of multiple mobile robots based on decentralized motion planning and negotiation," in *IEEE International Conference on Robotics and Automation*, 1997, pp. 3526–3533.
- [5] H. Chitsaz, J. O’Kane, and S. LaValle, "Exact pareto-optimal coordination of two translating polygonal robots on an acyclic roadmap," in *IEEE International Conference on Robotics and Automation*, vol. 4, 2004, pp. 3981 – 3986.

- [6] Z. Z. L. Chun and W. Chang, "A decentralized approach to the conflict-free motion planning for multiple mobile robots," in *IEEE International Conference on Robotics and Automation*, 1999, pp. 1544–1549.
- [7] J. O. J.P. Desai and V. Kumar, "Controlling formations of multiple mobile robots," in *IEEE International Conference on Robotics and Automation*, 1998, pp. 2864–2896.
- [8] H. Yamaguchi and J. Burdick, "Asymptotic stabilization of nonholonomic mobile robots forming group formations," in *IEEE International Conference on Robotics and Automation*, 1998, pp. 3573–3580.
- [9] S. M. R. C. M. Clark and J. Latombe, "Motion planning for multiple robots using dynamic networks," in *IEEE International Conference on Robotics and Automation*, 2003, pp. 4222–4227.
- [10] 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.
- [11] E. Klavins, "Communication complexity of multi-robot systems," in *Proc. Fifth International Workshop on the Algorithmic Foundations of Robotics*, Nice, France, 2002.
- [12] 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.
- [13] 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.
- [14] H. Tanner, A. Jadbabaie, and G. Pappas, "Flocking in fixed and switching networks," 2004, submitted: Transactions on Automatic Control.
- [15] R. Olfati-Saber, "Flocking for multi-agent dynamic systems: Algorithms and theory," *IEEE Trans. on Automatic Control*, vol. 51, no. 3, pp. 401–420, 2006.
- [16] S. Loizu, D. Dimarogonas, and K. Kyriakopoulos, "Decentralized feedback stabilization of multiple nonholonomic agents," in *IEEE International Conference on Robotics and Automation*, vol. 3, 2004, pp. 3012 – 3017.
- [17] 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.
- [18] R. Ghosh and C. J. Tomlin, "Maneuver design for multiple aircraft conflict resolution," in *Proc. American Control Conf.*, Chicago, IL, 2000.
- [19] C. Tomlin, I. Mitchell, and R. Ghosh, "Safety verification of conflict resolution maneuvers," *IEEE Trans. Intelligent Transportation Systems*, vol. 2, no. 2, 2001.
- [20] 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.
- [21] 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.
- [22] R. Tempo, G. Calafiore, and F. Dabbene, *Randomized Algorithms for Analysis and Control of Uncertain Systems*, ser. Communications and Control Engineering. Springer-Verlag, 2003.
- [23] 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.
- [24] M. S. Branicky, "Studies in hybrid systems: Modeling, analysis, and control," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, MA, 1995.
- [25] J. Lygeros, "Hierarchical hybrid control of large scale systems," Ph.D. dissertation, University of California, Berkeley, CA, 1996.
- [26] 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.
- [27] 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.
- [28] L. Pallottino, V. G. Scordio, E. Frazzoli, and A. Bicchi, "Probabilistic verification of a decentralized policy for conflict resolution in multi-agent systems," in *IEEE International Conference on Robotics and Automation*, 2006.
- [29] 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.