Probabilistic verification of decentralized multi-agent control strategies: a Case Study in Conflict Avoidance

Lucia Pallottino, Antonio Bicchi and Emilio Frazzoli

Abstract— **Many challenging verification problems arise from complex hybrid automata that model decentralized control systems. As an example, we will consider decentralized policies that steer multiple vehicles in a shared environment: properties of safety and liveness, such as collision avoidance and ultimate convergence of all vehicles to their goals, must be verified. To formally verify the behavior of proposed policies, it is desired to identify the broadest class of start and goal configurations, such that safety and liveness would be guaranteed. Simple conditions are proposed to identify such a class: ideally, a formal proof that such conditions are necessary and sufficient for safety and liveness is requested. Unfortunately, in decentralized control frameworks classical approaches are difficult to apply. Hence, probabilistic verification method can be applied to quantify the accuracy and the confidence of the veridicity of the desired predicate. The probabilistic verification method is applied to a recently proposed cooperative and completely decentralized collision avoidance policy for non-holonomic vehicles.**

I. INTRODUCTION

In recent years, Multi-Agent Systems (MASs) have attracted increasing attention and have been proposed for several applications, such as air traffic management, planetary exploration, and surveillance. MASs 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 [1]. 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 [2], [3], [4], [5], [6], [7], [8]. Although correct and complete centralized algorithms for the traffic management problem may exist, they typically require a large amount of computational resources. Furthermore, centralized approaches are typically liable to faults of the decision maker.

The emergence of reliable and economically viable wireless networking and distributed sensing is going to enable a revolution in this field towards decentralized approaches, by which each agent plans its own trajectory based only on information limited to neighboring agents. 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. Several authors have considered decentralized control of multiple mobile agents (see e.g.[9], [10], [11]). In [12], the authors propose a hybrid control architecture with parallel problem solving which guarantees collision avoidance. In [13] 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.

In a decentralized control framework, the formal verification of the safety and effectiveness of collision avoidance policies is a non trivial problem. In general this is true for many complex hybrid automata. In order to fully characterized the behavior of the system under the action of the proposed policy, it is desired to determine the largest set of initial and desired agent configurations under which safety and liveness properties are satisfied; in particular, we consider a set of simple algebraic conditions that are conjectured to identify such a class. Unfortunately, the formal verification that such conditions are necessary and sufficient to ensure safety and effectiveness (liveness) properties often appears to be overwhelmingly complex. Classical methods for convergence of agents towards their desired configuration are not easily applicable. A possible approach is then to 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 [14]. 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. The application of such methods allows a quantification on the efficiency and accuracy of the obtained experimental results. Furthermore, a relation on the number of experiments and the desired efficiency and accuracy is given.

In this paper, a probabilistic method is applied to the

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 authors and do not necessarily reflect the views of the supporting organizations.

L. Pallottino and A. Bicchi are with Faculty of Engineering, University of Pisa, 56100 Pisa, Italy `l.pallottino@ing.unipi.it`

E. Frazzoli is with the Laboratory for Information and Decision Systems, Massachusset Institute of Technology, Cambridge, MA 02139, USA, `frazzoli@mit.edu`.

verification of safety and liveness properties of decentralized conflict resolution polices of multi-agent systems. The paper is intended as a tutorial exposition on methods and theirs application, and expand upon material resented in [16]. A particular test case based on the Generalized Roundabout Policy described in [15] is considered.

II. HYBRID MODEL FOR DECENTRALIZED CONTROL

A crucial aspect in the development of decentralized policies for the networked multi-agent coordination problem is how the set of data used by each component of the system to elaborate solutions influences performance and safety requirements. In the case of a centralized system, the problem becomes trivial since all the data must be available to a single Decision Maker. Conversely, for networked decentralized systems, where data can be delivered through a limited bandwidth channel, a communication scheme "from all to all" may not be supported in time compatible with safety and performance. Therefore, the choice the *Information Structure* (IS), i.e. the set of data available at any time to each agent to decide its strategy, is a critical issue in networked mobility systems.

Information Structures are qualified in several categories. The *topology* of the IS describes the *relevance connectivity* of the communication network. In networked mobility, it appears natural to choose metric-based IS topologies, which model scenarios where each agent can share data only with agents in the distance range of their sensors/communication capabilities. In such schemes, relevance is thus a function of distance: typically, the ith agent has information regarding all other agents which are at a distance less than a given "Alert" radius A_i. In other terms, data concerning agent j are known to agent i only if the distance $d(i, j)$ is less than A_i, at a given refresh rate in time.

The size of the alert distance can then be regarded as a degree of centralization/decentralization. Indeed, very large alert distances relative to the maneuvering capabilities of the agents are tantamount to centralized control, as every agents gets full information on the system while still far away from conflicts. On the other hand, for small alert distances, a miopic resolution policy of one conflict might give raise to a cascade effect on other conflicts, with possibly destabilizing consequences.

Important topological aspects of the IS are *simmetry* and *transitivity*. Let $S_i(\tau)$ denote the set of indices of agents within distance A_i from the i–th agent at time τ. An information structure is *simmetric* if $i \in S_j \Rightarrow j \in S_i$; it is *transitive* if $i \in S_j$ and $j \in S_k \Rightarrow i \in S_k$. In other words, simmetry of an IS implies that if agent *A* has information about *B* then agent *B* has information about *A* (in formulae, $A_i = A_j \; \forall i, j$ in a metric-based topology). Transitivity implies that if agent *A* has information about *B* and agent *B* has information about *C*, then *A* has information about *C*. Some illustrative examples are reported in fig. 1. Notice that different alert distances cause non simmetry. Hence, simmetry implies $A_i = A_{alert}, \forall i$.

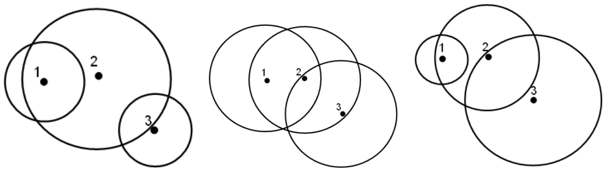

Fig. 1. Several different information structures for three vehicles. Left: $S_1 = \{1\}$, $S_2 = \{1, 2\}$, $S_3 = \{3\}$, non simmetric case . Middle, non transitive: $S_1 = \{1, 2\}$, $S_2 = \{1, 2, 3\}$, $S_3 = \{2, 3\}$; transitive: $S_1 = S_2 = S_3 = \{1, 2, 3\}$; Right, non transitive: $S_1 = \{1\}$, $S_2 = \{1, 2\}$, $S_3 = \{2, 3\}$; Right, transitive: $S_1 = \{1\}$, $S_2 = \{1, 2\}$, $S_3 = \{1, 2, 3\}$.

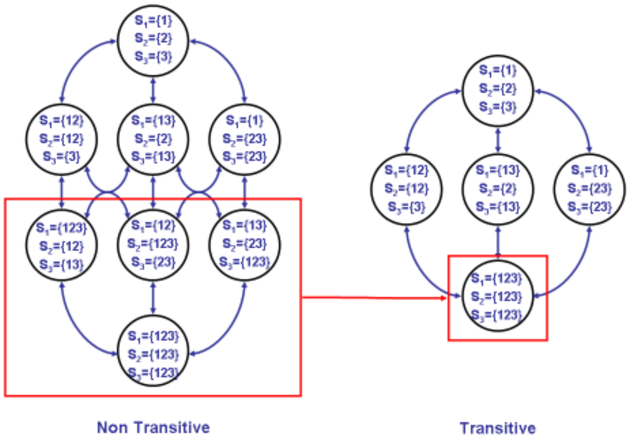

Fig. 2. Graphs describing the hybrid system whose nodes are different Information Structures for $N = 3$ agents. Arcs represent switching between different IS when an agent enters or exits the alert neighborhood of another agent.

Introducing transitivity may considerably simplify the analysis. In figure 2, two simmetric IS structures and their possible evolutions are represented for comparison. The graph on the right is relative to a transitive IS, which determines the collapse of some nodes into one equivalent node.

The information content (*semantics*) is another key issue defining the IS, with direct bearing upon the feasible refresh rates for closeby agents and on the number of supportable neighbors. Information on relevant agents may include their (relative) state (position, orientation, velocity), destination, and other factors. For simplicity, we will mainly distinguish two possible types of semantics: state–only, and state–destination messages. The difference is clearly in the fact that, while state–only information to an agent can be afforded with on-board sensors on the agent itself, destination information implies active transmission by each relevant agent.

Once the Information Structure has been determined, for the specific application problem, a collision avoidance policy can be designed. Obviously, the collision avoidance policy is based on the chosen IS for the level of the decentralization (dimension of A_i), the quantity of information exchanged by the agents (transitivity) and the type of information exchanged (semantics).

III. FUNDAMENTAL PROPERTIES OF THE COLLISION AVOIDANCE POLICIES

The big issue with decentralized schemes is obviously that switching among different modes can lead to situations where no feasible (in terms of safety and liveness) solution exist. The formal verification that under given condition a feasible solution exists is then a fundamental problem to deal with in case of decentralized control schemes.

We now introduce and formalize the fundamental properties a decentralized collision avoidance policy should verify.

A. Admissibility

In general, the total number of vehicles in the system may vary in time. Hence, we 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 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:

$\mathbf{P_1}$: A configuration set $G = \{g_i,\ i = 1, \ldots, n\}$, is *unsafe* for the policy π if there exists a set of target configurations $G_f = \{g_{f,i},\ i = 1, \ldots, n\}$ such that application of π leads to a collision;

$\mathbf{P_2}$: A target configuration set $G_f = \{g_{f,i},\ i = 1, \ldots, n\}$, is *blocking* for the policy π if there exists a set of configurations $G = \{g_i,\ i = 1, \ldots, n\}$ 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 $\neg\mathbf{P}_1\left(G(t)\right) \wedge \neg\mathbf{P}_2(G_f)$.

Clearly, tests to check the two properties $\mathbf{P_1}$ and $\mathbf{P_2}$ are needed for each control policy.

B. Safety

For a generic policy π we need a test $\mathbf{T}_1(\pi, G)$ based on initial configuration such that if $\mathbf{T}_1(\pi, G)$ is verified the property $\neg\mathbf{P_1}(\mathbf{G})$ is implied.

In general for "safety" it is meant that any two agents maintain a minimum given distance during the motion. The proposed test-case, i.e. the Generalized Roundabout Policy (GRP), has been developed for agents with kinematics

$$\left\{\begin{array}{lcl} \dot{x}_i(t) & = & v_i\cos(\theta_i(t)) \\ \dot{y}_i(t) & = & v_i\sin(\theta_i(t)) \\ \dot{\theta}_i(t) & = & \omega_i(t) \end{array}\right. \qquad (1)$$

where $\omega_i : \mathbb{R} \rightarrow [-\frac{v_i}{R_C}, \frac{v_i}{R_C}]$ is a bounded signed curvature control signal and v_i is a constant linear velocity. In this framework a *reserved region*, over which each active agent claims exclusive ownership is defined as follows. 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 $\omega = -\frac{v_i}{R_C}$. In other words,

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

A test $\mathbf{T}_1(\pi, G)$ for property $\mathbf{P_1}$ in case of the GRP is provided by the following theorem.

Theorem 1: If the reserved disks of at least two agents in G overlap, property $\mathbf{P_1(G)}$ is verified. In other words, if $\|c(g_i) - c(g_j)\| \geq 2 + d_s$, $\forall i, j \in \{1, \ldots, n\}, i \neq j$, it follows that $\forall t \geq 0$, $d(g_i(t), g_j(t)) > d_s$, $\forall \in \{1, \ldots, n\}, j \neq i$.

C. Liveness

The analysis of property $\mathbf{P_2}$ is in general complex, and testable conditions are needed also for this property. Hence, for a generic policy π we also need a test $\mathbf{T}_2(\pi, G_f)$ based on final configuration such that it implies $\neg\mathbf{P_2}(\mathbf{G_f})$.

In case of the GRP, property $\mathbf{P_2}$ hinges upon the definition of a condition concerning the separation of reserved discs associated with target configurations. A sparsity condition on the target configuration has been introduced: any circle of a given radius (that depends on the dimension of the safety disk and on $0 \leq m \leq n$), can contain at most $m - 1$ reserved disk centers of targets. In [16] a formal definition of the sparsity condition may be found.

Consider the property:
$\mathbf{P_3}$: A target configuration set $G_f = \{g_{f,i},\ i = 1, \ldots, n\}$ is *clustered* if the sparsity condition is violated.

In case of the GRP we are able to prove the following theorem

Theorem 2 (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. $\mathbf{P}_3(G_f) \Rightarrow \mathbf{P}_2(G_f)$.

Concluding, for a control policy π conditions that are necessary or sufficient for the safety and liveness of the overall system must be determined. In other word, tests $\mathbf{T}_1(\pi, G)$ and $\mathbf{T}_2(\pi, G_f)$ such that $\mathbf{T}_1(\pi, G) \wedge \mathbf{T}_2(\pi, G_f) \Rightarrow \neg\mathbf{P}_1\left(G(t)\right) \wedge \neg\mathbf{P}_2(G_f)$ or such that $\neg\mathbf{P}_1\left(G(t)\right) \wedge \neg\mathbf{P}_2(G_f) \Rightarrow \mathbf{T}_1(\pi, G) \wedge \mathbf{T}_2(\pi, G_f)$ are needed.

Let us suppose that, as in the GRP case, necessary conditions have been determined. Often, the formal verification that such conditions are sufficient to ensure safety and liveness appears to be overwhelmingly complex. In next section, a probabilistic approach to the problem sufficiency verification is described.

IV. PROBABILISTIC VERIFICATION OF A DECENTRALIZED CONTROL POLICY

We briefly introduce the probabilistic method for reader convenience (for more details, see e.g. [14]).

Given a dynamical system D subject to uncertainties Δ and a predicate P_D defined on D which we want to verify. Let $\mathcal{B}$ be the bounded set in which uncertainties are confined and $f_\Delta(\Delta)$ the associated probability density function. Probabilistic verification consists in evaluating, with a prescribed confidence, the probability

$$p_D := \mathrm{PR}_{\mathbf{\Delta}}\{P_D(\mathbf{\Delta})\} = \int_{\mathcal{G}} f_{\mathbf{\Delta}}(\Delta)d\Delta,$$

where $\mathcal{G} \subseteq \mathcal{B}$ denotes the *good set* of $\Delta \in \mathcal{B}$ for which $P_D(\Delta) = true$.

Given a *performance function* $J_D(\mathbf{\Delta})$ of system D, the probability that a given performance level γ is attained under uncertainties as above can be expressed by the predicate $P_D(\mathbf{\Delta}) = \{J_D(\Delta) \leq \gamma\}$.

The measure of the predicate veridicity is given by the volume ratio $r = Vol(\mathcal{G})/Vol(\mathcal{B})$ that can be evaluated by a Monte Carlo approach if a uniform distribution function on $\mathcal{B}_d$ is considered. Indeed, under this assumption $p_D = r$. Let us denote by $\mathbf{\Delta}^i$, $i = 1, \ldots, N$ N random samples within $\mathcal{B}$. An estimate of r based on the empirical outcomes of the N instances of the problem is given by $\hat{p}_D(N) = \frac{1}{N}\sum_{i=1}^{N} I_{\mathcal{G}}(\mathbf{\Delta}^i)$ where $I_{\mathcal{G}}(\mathbf{\Delta}^i) = 1$ if $\mathbf{\Delta}^i \in \mathcal{G}$ and 0 otherwise.

This result provides a finite N such that the empirical mean $\hat{p}_D(N)$ differs from the true probability p_D less than ϵ with probability greater than $1-\delta$, i.e. $Pr\{|p_D - \hat{p}_D(N)| < \epsilon\} > 1-\delta$, for $0 < \epsilon, \delta < 1$. To determine the minimum number N the Chernoff bound [17] can be used:

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

Notice that the sample size N, given by (2), is independent of the size of $\mathcal{B}$ and of the distribution $f_{\mathbf{\Delta}}(\Delta)$.

A. Application of the probabilistic method

Consider the following statement:

Conjecture [Sufficient conditions for admissibility] The policy π provides a safe and non-blocking solution for all plans (G_0, G_f) that verify $\neg\mathbf{T}_1(\pi) \wedge \neg\mathbf{T}_2(\pi)$.

For the GRP policy the conjecture is: the GR policy provides a non-blocking solution for all safe and non clustered plans (G_0, G_f).

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

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

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

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

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.

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 $\mathcal{G} = \{(G_0, G_f) \in \mathcal{B} | \mathbf{P}'_\pi(G_0, G_f)\}$ denote the "good" set of problem data for which the predicate applies. Also, let $\mathcal{C} = \{(G_0, G_f) \in \mathcal{B} | \neg\mathbf{T}_1(\pi) \wedge \neg\mathbf{T}_2(\pi)\}$ denote the set of safe and non clustered plans.

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

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

can be regarded as a measure of the probability of correctness of the conjecture. As reported above, a number N of experiments must be conducted in order to have that empirical mean $\hat{r}(N)$ differs from the true probability r less than ϵ with probability greater than $1-\delta$.

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.

For the application of the probabilistic approach to the specific GRP refer to [16]. For the GRP, several experiments have been conducted to assess how these two indicators vary with the parameters (see Fig. 3 and 4). 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 5. From

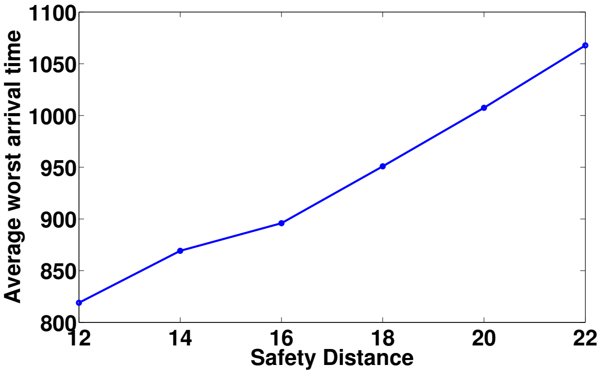

Fig. 3. Average worst arrival time (over 300 experiments) vs. safety distance, for a system of 10 agents. The average unconstrained solution time is close to 520.

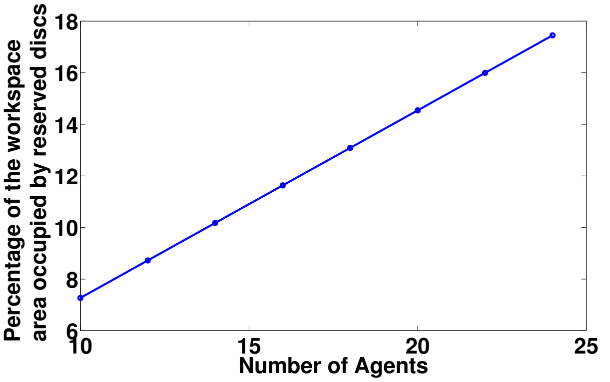

Fig. 4. Percentage of workspace area occupied by agents and their reserved discs for different numbers of agents.

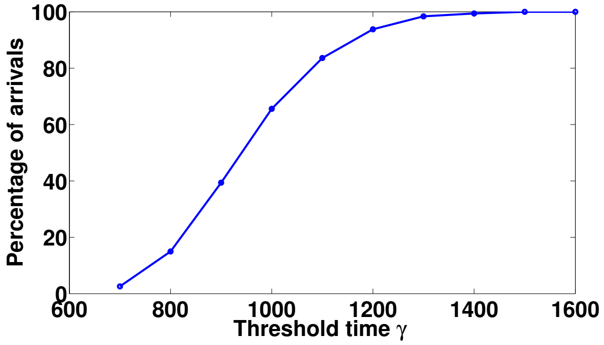

Fig. 5. Percentage of arrivals with respect to threshold time γ.

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 (2) 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 $d_s = 18$.

B. *Qualitative evaluation of the condition and of the liveness of the policy*

It may be interesting to provide a qualitative evaluations of tests $\mathbf{T}_1(\pi, G)$ e $\mathbf{T}_2(\pi, G_f)$.

The dimension of $\mathcal{C}$ in $\mathcal{B}$ usually depends on the value of the number of agents n and the value of the associated safety radius R_S. Figure 6 represents the normalized dimension of $\mathcal{C}$ in $\mathcal{B}_n$ with respect to variation of $n \in \{2, \dots, 20\}$ and $R_S \in \{2, \dots, 40\}$ for the GRP. In figure 7 the z-axis view is reported. In case of the GRP, projections of the isodimensional curves on the $(n,\, R_S)$ plane appear to be hyperbolas, i.e. $n\, R_S = const.$.

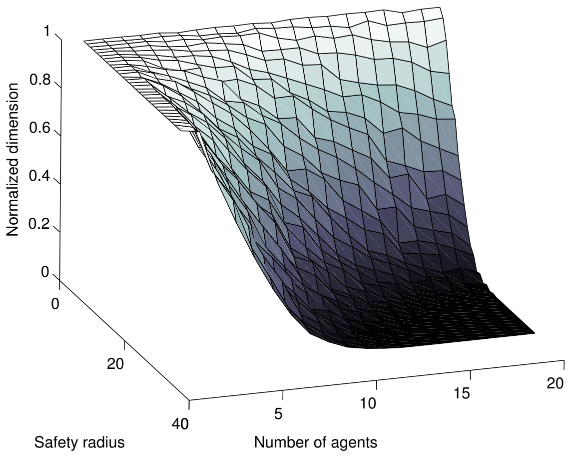

Fig. 6. The normalized dimension of $\mathcal{C}$ in $\mathcal{B}$ with respect to variation of n and R_S.

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 GRP 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

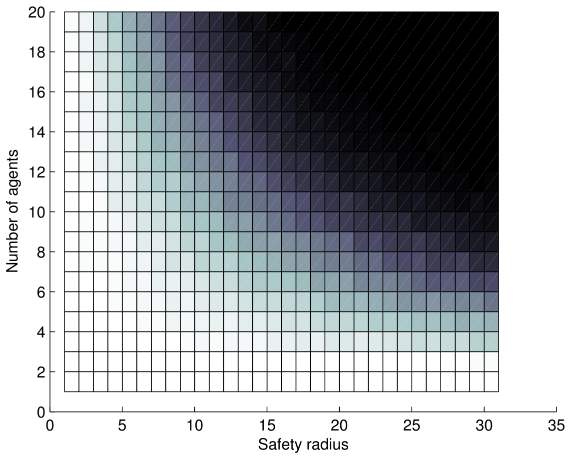

Fig. 7. Projections of the isodimensional curves on the $(n,\, R_S)$ plane appear to be hyperbolas.

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.

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 $7\text{meter} \times 8\text{meter}$ we are able to manage safely 10 agents with a safety disc diameter of 60 centimeters.

A similar approach can be used to qualitatively evaluate tests $\mathbf{T}_1(\pi, G)$ e $\mathbf{T}_2(\pi, G_f)$ for a generic decentralized control policy π.

V. CONCLUSION

The probabilistic method proposed in [14] has been applied to probabilistically verification conditions for safety and liveness of decentralized control policies for collision avoidance. Indeed, decentralized control problems are modeled with complex hybrid automata from which challenging verification problems arise. For each decentralized control policies some testable conditions on start and goal configurations are needed. Then, a formal proof that such conditions are necessary and sufficient for the safety and the liveness of the overall system would be required. Unfortunately, as e.g. in the GR policy only one of the two implication may be analytically proved. For the other implication (in the GRP the sufficiency implication) a probabilistic approach may be used. The application of such methods allows a quantification on the efficiency and accuracy of the obtained experimental results. Furthermore, a relation on the number of experiments to be done and the desired efficiency and accuracy is given.

REFERENCES

[1] Nikos Vlassis, "A Concise Introduction to Multiagent Systems and Distributed AI", downloadable from http://carol.science.uva.nl/ vlassis/cimasdai/, University of Amsterdam 2003.

[2] LaValle, S.M., Hutchinson, S.A.: Optimal motion planning for multiple robots having independent goals. IEEE Trans. on Robotics and Automation **14** (1998) 912–925

[3] Gerkey, B.P., Matariĉ, M.J.: Sold!: Auction methods for multi-robot coordination. IEEE Trans. Robotics and Automation **18** (2002) 758–768

[4] C. Tomlin, G. J. Pappas, and S. Sastry, "Conflict resolution for air traffic management: A case study in multi-agent hybrid systems," *IEEE Trans. on Aut. Contr., vol. 43, no. 4, pp. 509–521, April 1998.*

[5] P. K. Menon, G. D. Sweriduk, and B. Sridhar, "Optimal strategies for free-flight air traffic conflict resolution," *Journ. of Guidance, Control, and Dynamics, vol. 22, no. 2, pp. 202–211, March–April 1999.*

[6] E. Frazzoli, Z.H. Mao, J.H. Oh and E. Feron, "Resolution of Conflicts Involving Many Aircraft via Semidefinite Programming" to appear *AIAA, Journal of Guidance, Control and Dynamics, 2001.*

[7] A. Bicchi and L. Pallottino, "On Optimal Cooperative Conflict Resolution for Air Traffic Management Systems" *IEEE Transactions On Intelligent Transportation Systems, vol. 1, no.4, pp.221-231, Dec. 2000.*

[8] L. Pallottino, E. Feron and A. Bicchi, "Conflict Resolution Problems for Air Traffic Management Systems Solved with Mixed Integer Programming" *IEEE Transactions On Intelligent Transportation Systems, vol. 3, no.1, pp.3-11, March 2002.*

[9] Lumelsky, V.J., Harinarayan, K.R.: Decentralized motion planning for multiple mobile robots: the cocktail party model. Autonomous Robots **4** (1997) 121–35

[10] Klavins, E.: Communication complexity of multi-robot systems. In: Proc. Fifth International Workshop on the Algorithmic Foundations of Robotics, Nice, France (2002)

[11] Inalhan, G., Stipanovic, D.M., Tomlin, C.J.: Decentralized optimization, with application to multiple aircraft coordination. In: Proc. IEEE Conf. on Decision and Control. (2002)

[12] K. Azarm and G. Schmidt, "Conflict-Free Motion of Multiple Mobile Robots Based on Decentralized Motion Planning and Negotiation", *IEEE Int. Conf. on Robotics and Automation*, pp 3526-3533, 1997.

[13] C. M. Clark, S. M. Rock and J.C. Latombe, "Motion Planning for Multiple Robots using Dynamic Networks", Proc. Int. Conf. on Rob.and Autom, pag. 4222-4227, 2003.

[14] R. Tempo, G. Calafiore, and F. Dabbene, *Randomized Algorithms for Analysis and Control of Uncertain Systems*, ser. Communications and Control Engineering. Springer-Verlag, 2003.

[15] L. Pallottino, V. G. Scordio, E. Frazzoli, and A. Bicchi, "Decentralized Cooperative Conflict Resolution for Multiple Nonholonomic Vehicles." In Proc. AIAA Conf. on Guidance, Navigation and Control, 2005.

[16] 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.

[17] H. Chernoff, "A measure of asymptotic efficency for tests of a hypotesis based on the sum of observations," *Annals of Mathematical Statistics*, vol. 23, pp. 493–507, 1952.