Behavioral control of robot teams

João Sequeira

Institute for Systems and Robotics
Instituto Superior Técnico
Lisbon, Portugal
Email: jseq@isr.ist.utl.pt

Isabel Ribeiro

Institute for Systems and Robotics
Instituto Superior Técnico
Lisbon, Portugal
Email: mir@isr.ist.utl.pt

Abstract

This paper discusses the modeling and control of fully autonomous and semi-autonomous robots, operating individually or as teams, in a behavioral context.

The recent developments in networking technologies foster the use of teams of robots in a variety of applications, e.g., surveillance and rescueing. Additionally, the inclusion of human agents in a team of robots raises relevant questions on the modeling of human-robot interactions. To fully explore the potential of robot teams these interactions must be accounted for at multiple stages in the design of a robot control architecture.

The paper presents a behavioral paradigm for the control of robots supported on basic concepts from nonsmooth calculus. The paradigm develops into a standard two layer hybrid architecture encompassing basic motion primitives and supervision. This approach allows the representation of semantic concepts, a key issue in human-robot interactions, in a natural way by shaping goal regions for the robots to reach in their configuration spaces. Robot teams are handled by extending the supervision layer, designed for each robot, to account for the interactions among teammates.

I. Introduction

The recent developments in computing and communications technologies is pushing robotics through a transformation period. Despite the intensive research in the numerous areas related with robotics, fully autonomous robots (FAR) are not yet available to execute even moderately complex missions in realistic scenarios. Robustness to uncertainties and intelligence, considered as the rational decision making capability, are still too weak for a robot to be left taking decisions completely on its own. The recent interest in Mars exploration pushed the development of control architectures towards semi-autonomous robotic (SAR) systems. This is an example where human supervision is paramount to handle, for instance, contingency scenarios.

Robot teams are the natural extension to single robot systems and currently a key area in robotics research. In the realm of the growing interest in this area are the aforementioned FAR systems weaknesses, namely when operating isolatedly, and the fact that the performance of teams of humans/animals is in general greater than when the individuals operate isolatedly. With robot teams, working cooperatively, a wider variety of missions can be executed, and the use of sensing resources can be optimised.

Networked systems, of which robot teams are a subclass, have been studied in the context of dynamical systems (see for instance the survey in [1]). The modeling and control issues in classical robotics are nowadays well understood (though nonlinear systems is still an active research field). Cooperative robotics is transversal to multiple scientific areas and, in general, it is possible to identify in each of them different abstraction levels. When the control of the team is considered at a high level of abstraction, as when the communication among robots is of behavioral/linguistic type, the analysis of the system requires mathematical tools to handle the semiotics of the interactions. This paper proposes the use of basic concepts of nonsmooth systems, developed into a behavioral approach, to the control of robots operating individually or within a team.

Behavior based robotics has been a hot topic in the last two decades motivated by the fact that human/animal behavior is a key factor in the success of the species. The early behavior-based paradigms essentially account for the robustness issue, with intelligence being regarded as a byproduct of the complex interactions among primitive behaviors. Models of personality developed since the 50's by multiple researchers in Psichology, [14], can easily led to a behavior definition. However, its usefulness in the formal analysis of the robot control architectures was reduced. The framework developed in [24], supported on dynamical systems theory, gave a sound support to the behavior notion. However, it failed to capture some of the subtleties that characterize human behaviors, namely the pragmatics of the communications between humans. Despite these limitations, behavioral concepts were introduced into classical models up to a point where some behavior concept can be identified in almost any architecture, [3], [7], [8], [9], [12], [21].

Full autonomy has been, since the early days, the main goal in the design of control architectures, with semi-autonomy being given a secondary role (in what concerns scientific research). However, semi-autonomy offers the additional challenge of having teammates (the human operators) that are potentially unaware of, for instance, kinematic motion constraints, limited communications bandwith or limited linguistic understanding capabilities, and yet try to interact with the robots. Though multiple SAR control architectures, [2], [6], [17], [23], have been proposed along the years, most of them sharing conceptual relationships with FAR control architectures (e.g., behavioral concepts), there is still need for research on the modeling of such human-robot interactions which may yield improved communication schemes for FAR systems.

When a human joins a team of robots, forming a SAR team, the interaction human-robot becomes a key issue. The native forms of communication among humans are still too complex for robots to mimic. However, their richness is inspirational to the modeling of human-robot and robot-robot interactions. Communication protocols, environment perception and learning/adaptation strategies are only a few examples of the related key research areas in SAR control paradigms. These topics have been traditionally considered within specific frameworks (for instance, semiotics problems have been considered within AI techniques and learning has been considered both in control theory and in AI areas). Whether it is done under graphics, voice or command line based environments, the interface between humans and robots translates the commands specified in the operators language to motion commands and hence it affects most of the building blocks in a control architecture.

For SAR teams, it is of uttermost importance to design control architectures providing simple/intuitive interaction between humans and robots. One of the aims of this research is to motivate for the use of basic concepts in nonsmooth systems that provide intuitive tools to design such interfaces and to model SAR systems. Though the potential to handle semantics, the role of nonsmooth systems in the modeling of the multiple semiotic aspects (e.g., pragmatics and syntactics) is open for discussion and will not be addressed in this paper.

The paper is organised in three sections. Section II describes a behavioral model for a robot control architecture, with emphasis on SAR applications, and its extension to the control of teams of robots. Simple, though relevant, simulation results illustrate the operation of the architecture. Section III summarizes the relevant topics identified along the paper.

II. Behavioral modeling of semi-autonomous robot systems

The natural differences in the interpretation of a mission by different users, i.e., the ambiguity in mission interpretation, may lead to different commands being sent to the robot leading to different mission executions, eventually successfull all of them. Handling language ambiguities is an issue in semiotics. In this paper ambiguity is considered in the perspective of semantics. Dealing with ambiguities in this perspective has been in the foundations of the traditional behavioral approaches but can also be found in classical hierarchical architectures under the principle of increasing intelligence with decreasing precision, [16]. The framework described in this paper explores this characteristic of having different commands resulting in a similar behavior.

Identifying intelligence with the ability to define a behavior (with the proper degree of uncertainty accounted for) leads to the conclusion that (i) if robots are to be intelligent, formal behavior definitions are needed, and (ii) these must be of low complexity, as precision may be relaxed when specifying a mission. For instance, if a human is specifying a mission for either a FAR or SAR the precision may be decreased simply by relaxing the mission definition. Instead of reaching a specific configuration or following a specific trajectory, a mission may aim at reaching a region or tracking a sequence of regions in the configuration space. The practical consequence of this reasoning is that by specifying goal regions (instead of specific goal configurations) a human or a robot can account for semantics in the information exchanged with teammates without compromising the mission objectives.

It is along these lines that the paper proposes a control architecture with mathematical objects tailored to account for the aforementioned semantic relationships. The motion of a robot is seen as resulting from the application of a set of operators on a space of motion strategies defined to cope with the ambiguities in the mission interpretation.

In general, a robot can successfully execute a mission by choosing one path to follow among a set of feasible paths. Without loosing generality, one can assume that this set of feasible paths is entirely contained in a compact region of the robot configuration space. The resulting mathematical object, composed by this region (the action bounding region) and by some algorithmic process to generate the paths, is called an action. Actions can be thought of as motion strategies, motions trends or behaviors, and in a sense they represent the ability of the robot to move inside the action bounding region. By shaping this bounding region it is possible to control the ambiguity in the paths generated by an action and account, for instance, for the obstacles present in the environment and other positional mission constraints.

In a mission, the robot executes a sequence of actions. The actions in this sequence are composed through a special mathematical composition operator. Additional operators allow the transformation of an action into an equivalent one and the expansion of the action bounding region (see [19] for details). The result of the application of these operators to the space of actions is a path along which the robot is driven during the mission.

Under the aforementioned principles, the robot behavioral model is intrinsically hybrid as it involves discrete decision making over a set of continuous actions. Figure 1 represents such model where the symbols $\otimes$, $\boxtimes$

and ⊜ stand for the above referred action operators, namely action composition, action expansion and action equivalence transformation, respectively. At each event a supervisor, implemented as a finite state automaton, chooses the action to execute.

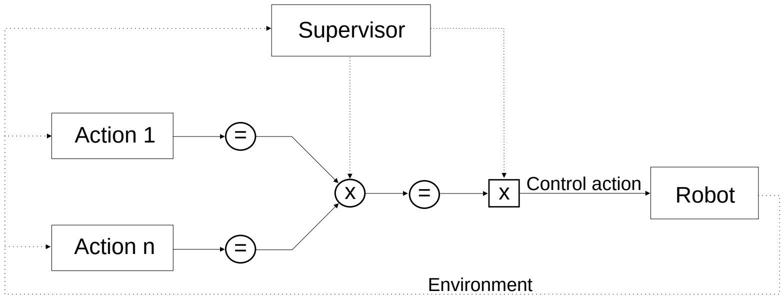

Fig. 1. Hybrid control architecture for a single robot

The initial conditions for an action do not need to coincide with the terminal conditions of the previous action. In this case a path to link the two consecutive actions is generated by the composition operator. Otherwise, the composition operator is simply the identity operator.

Formally, an action can be described by (i) its bounding region, accounting for positional mission and environmental constraints and (ii) a differential inclusion $\dot{q} \in F(q, u)$ where q (a solution trajectory of the inclusion, i.e., a sequence of configurations along time) and u stand, respectively, for the path to follow and the control variables and $F(\cdot)$ stands for the set of admissible velocities at each q. F represents the intersection between the set of feasible velocities (resulting from the robot kinematics/dynamics) and a set of admissible velocities defined by the mission. Computing a solution to this differential inclusion constrained to have q inside the action bounding region is a problem in viability theory that has been treated extensively in the literature (see for instance [4], [5]). When F is a singleton the differential inclusion reduces to an ordinary differential equation and the actions may be seen as ordinary algorithms that produce a unique reference path for the robot to follow.

The design of the supervisor is very much dependent on the mission and on the environment. However, the intrinsic properties of the architecture in Figure 1 can be used to set the design guidelines for the supervisor. A key result in nonsmooth calculus is given by the generalised Lyapunov stability theorem (see for instance [22]). If a mission goal is identified with an equilibrium state of the overall system, the referred theorem roughly states that if there exists a function (the Lyapunov function) of a performance index upper bounded by a strictly decreasing continuous function then the assymptotic stability of the equilibrium state is guaranteed. This amounts to say that the mission is successfully completed.

Let λ stands for the referred performance index. In what concerns the mission objectives, the quality of the execution of an action is measured by this index. Without lack of generality each value of λ can be classified into one of three classes according to $\lambda > 0$, $\lambda = 0$ and $\lambda < 0$. Positive values indicate that the current action is performing acceptably, i.e., the robot is heading towards the mission goal. Null values indicate that the current action is neither moving the robot away from the mission goal nor is making it approaching the goal. Negative values of λ indicate that the current action is moving the robot away from the mission goal.

At each event, each class of the performance index λ is identified with a subset of the actions creating a partition $A = A_1 \cup A_2 \cup A_3$ in the space of the actions. The supervisor selects the action to execute based on the corresponding λ value. Table I presents a finite state automaton structure for the supervisor.

Step	Description
1	compute the partition $A_1 \cup A_2 \cup A_3$ of the set of actions
2	if $A_1 \neq \emptyset$ choose any action in A_1
3	else, if $A_2 \neq \emptyset$ choose any action in A_2
4	else, choose an action in A_3.

TABLE I

Finite state automaton structure of the supervisor

The control architecture for a robot operating within a team, either containing only robots or including humans, remains unchanged. The locomotion capabilities of a robot when operating within a team stay basically the same as when it operates individually and hence its set of actions remains unaltered. In general, accounting for the interactions between robots requires additional states in the supervisor, expressing how the actions change in response to the interaction with the teammates. These states are embedded in the first step of Table I.

The key issue in what concerns SAR control lies in the interfacing between humans and robots. In general, the human-robot interaction (HRI) goes beyhond the human-computer interaction as it accounts for the so called "human factors", [11]. Among these factors, one can include the semantic differences that motivate this behavioral approach, together with other complementary semiotic aspects. Furthermore, issues such as the anthropomorphic characteristics of the robots become important when designing the interface as humans tend to interact better with human-like robots, [11].

Figure 2 illustrates the extension of the basic model in Figure 1 for SARs. The relevant characteristic in this model is the inclusion, in the supervisor block, of a HRI interface and of a the layer of abstractions to account for semantics. Abstractions are sets of actions built uppon different configurations of the basic locomotion capabilities. For instance, a "move forward fast" is an abstraction for "move forward with a specific velocity". A set of "move ... fast" commands represents a "fast moving robot" abstraction. Such abstractions tend to simplify both the robot-robot and the human-robot interactions as less information is exchanged between teammates.

A relevant subclass of abstractions is given by roles. A human may interact with a robot under a variety of roles for instance, as a "supervisor", "operator", "mechanic", "peer" or "bystander", [18]. Each of these roles grants specific priviledges to the human and configures the robot to react accordingly. In some sense, each role may be seen as assigning a particular semantic to each of the actions. For instance, a "move forward" action may have different limit velocities whether it is executed under the "operator" or the "mechanic" role.

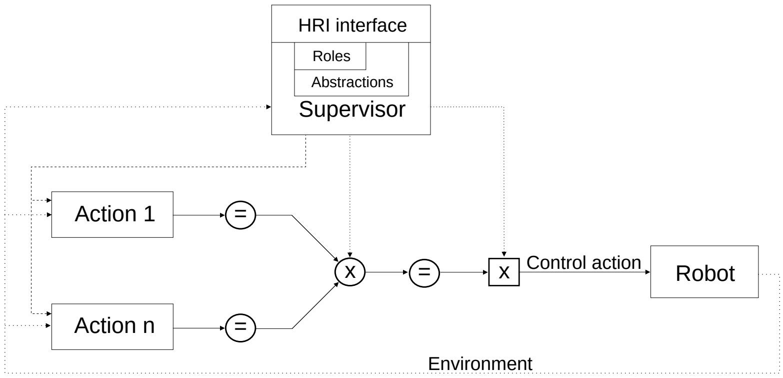

Fig. 2. Conceptual model of a semi-autonomous robot in a team

If the humans are assumed to have enough knowledge on the robots (operating isolatedly or in a team) the HRI can communicate with the human operators using standard computer languages. The widespreading of robot applications tends to make this assumption fail as an increasing number of non knowledgeable humans interact with the robots. Therefore, a key feature required to a HRI language must be the ability to cope with semantics (as mentioned in Section I this paper does not address the whole semiotics of a language). The ambiguities in human-human interactions amount to say that different language constructs are interpreted equivalently, that is, as synonyms. Establishing equivalence relationships, e.g., synonyms, between concepts requires the definition of a space of concepts, with adequate properties, beforehand. The architecture in Figure 2 implements this capability at the level of the actions. At the supervisor level the abstractions partially fulfil this objective.

Multiple computer languages have been used in robotics. Imperative languages, such as C++ and declarative languages like Haskell, [15] and FROB, [10], supported on Haskell, have been used in robot control. The languages explicitly designed for HRI must account for low complexity programming, communications and knowledge representation. RoboML, [13], supported on XML technology, is an example of such languages, though with limited semantic capabilities.

Additional details on the architecture in Figure 2, namely details on the mathematical operators, can be found in [20], [21]. The development of the upper semantic layer is ongoing work.

The following example illustrates the ability of a team of robots to move through a sequence of goals specified as arbitrary regions in the configuration space. The actions are obtained by modeling the velocity sets after cones obtained from the difference between the goal sets and the current configuration of the robot. This experiment aims primarily at showing the performance of a team of robots, each controlled under the architecture in Figure 1, competing to execute a mission loosely specified through the goal sets (see [20], [21] for additional details on the choice of the performance index and the supervisor finite state automaton implementation in similar experiments).

In this example a team of three unicycle robots is assigned the mission of reaching a sequence of three goal regions. The goal regions, shown as shadowed areas in Figure 3-a), were chosen (using a simple graphical interface for HRI) to have the robots interacting with each other to maintain a loose formation. The robots have no physical dimensions and each of them is equipped with three actions: pure rotation, pure translation and circular forward/backward motion. Figures 3-a) and 3-b) show the trajectories obtained, respectively in the configuration space and in the xy space. Figure 3-c) illustrates the evolution of the performance indices along the mission where it is seen the abrupt changes that correspond to the switching of actions. Figure 3-d) shows the minimal distance between each robot and its teammates. The supervisors allow the robots to be arbitrarily close to any teammate but not to collide with them.

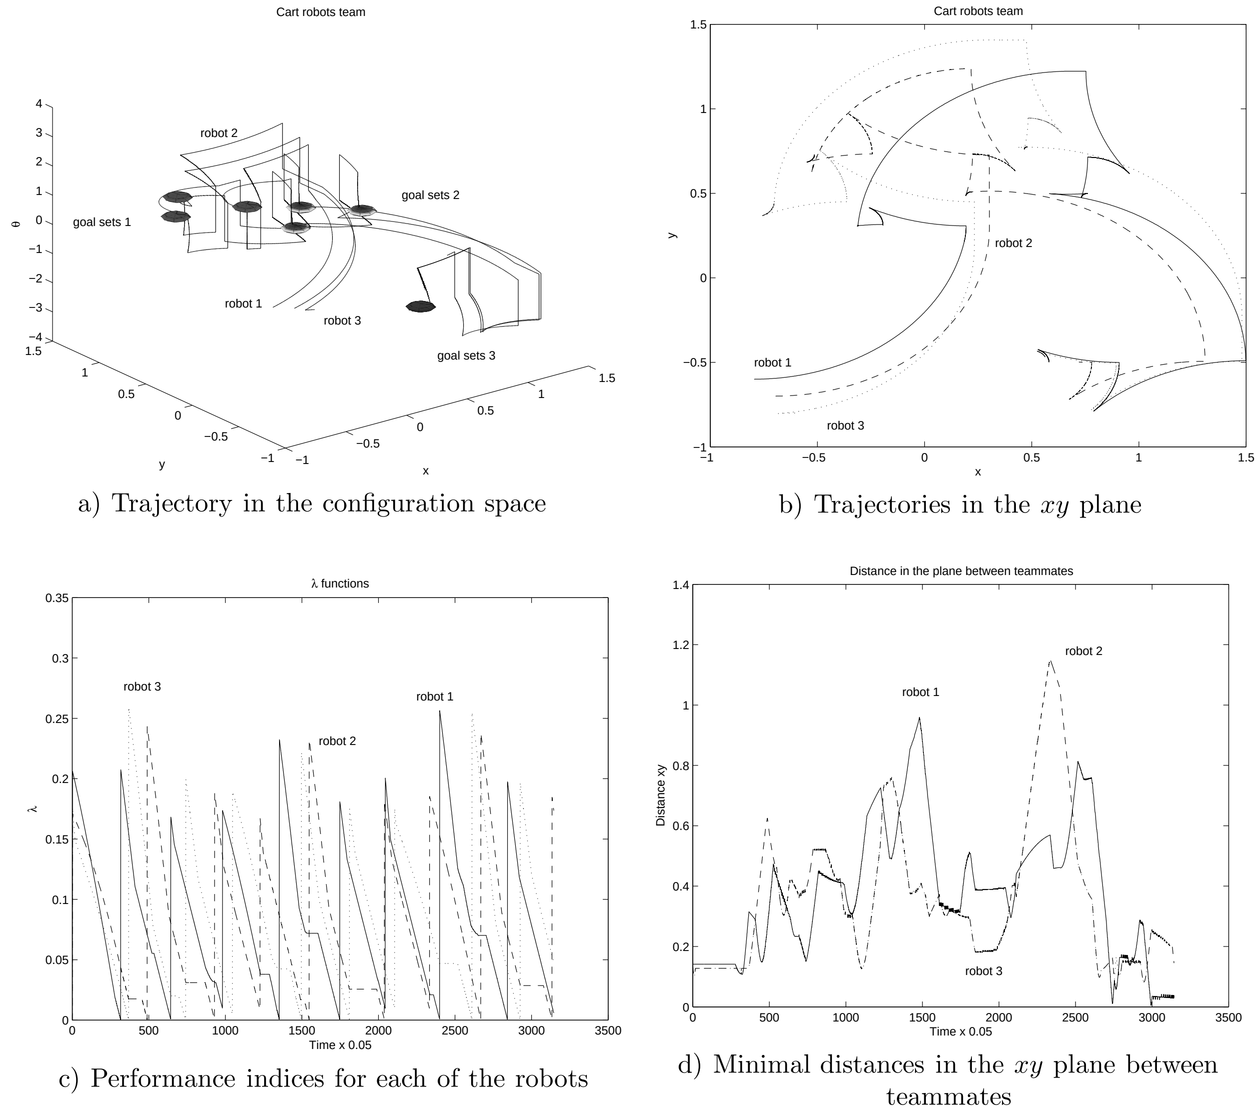

a) Trajectory in the configuration space

b) Trajectories in the xy plane

c) Performance indices for each of the robots

d) Minimal distances in the xy plane between teammates

Fig. 3. A team of 3 unicycle robots moving in a formation

III. Conclusions

The paper describes a behavioral hybrid architecture for robot teams (either SAR or FAR) supported on simple mathematical objects chosen to account with a rough model of semantic relationships in a mission.

The simple simulation results presented illustrate the main idea behind the architecture, namely that the complexity of robot control is largely reduced if the mission goals are loosely specified, i.e., specified up to a semantic transformation such as equivalence.

As a final remark, SAR systems are a challenging area to develop new control paradigms from which (almost paradoxical) FAR control can certainly profit. Additional research in human/animal interactions/behaviors is needed to allow a better characterization of the semiotic processes/requirements involved and their relevance to robotics. In parallel, alternative theoretical frameworks must be scanned for features that can match such requirements.

Acknowledgments

This work was supported by the FCT project SACOR - POSI/SRI/40999/2001 and Programa Operacional Sociedade de Informação (POSI) in the frame of QCA III.

References

[1] IEEE Control Systems Magazine - Complex Networks Part II. 22(1), February 2002.
[2] J. Albus. A control system architecture for intelligent systems. In Procs. of the 1987 IEEE Intl. Conf. on Systems, Man and Cybernetics, October, 20-23 1987. Alexandria, VA.
[3] J. S. Albus and A. M. Meystel. A Reference Model Architecture for Design and Implementation of Intelligent Control in Large and Complex Systems. International Journal of Intelligent Control and Systems, 1(1):15–30, 1996.
[4] J. P. Aubin. Viability Theory. Birkhäuser, 1991.
[5] J.P. Aubin and A. Cellina. Differential Inclusions. Springer-Verlag, 1984.
[6] R. Aylett and D. Barnes. A multi-robot architecture for planetary rovers. In Procs. of the 5th European Space Agency Workshop on Advanced Space Technologies for Robotics & Automation, December 1998. ESTEC, The Netherlands.
[7] T. Balch and R. Arkin. Behavior-based formation control for multi-robot teams. IEEE Trans. on Robotics and Automation, 14:926–939, December 1999.
[8] E. Bicho and G. Schöner. The dynamic approach to autonomous robotics demonstrated on a low-level platform. In Procs. of the 4th Symposium on Intelligent Robotic Systems, SIRS'96, July 1996. Lisbon, Portugal.
[9] R. Brooks. A Robust Layered Control System for a Mobile Robot. IEEE Journal of Robotics and Automation, RA-2(1), 1986.
[10] G. D. Hager and J. Peterson. Frob: A transformational approach to the design of robot software. In Procs. of the 9th Int. Symp. of Robotics Research, ISRR'99, 1999. Snowbird, Utah, USA, October 9-12.
[11] S. Kiesler and P. Hinds. Introduction to the special issue on human-robot interaction. Human-Computer Interaction, 19(1-2), 2004.
[12] J. Košecká, H. Christensen, and R. Bajcsy. Experiments in behavior composition. Robotics and Autonomous Systems, 19:287–298, 1997.
[13] M. Makatchev and S.K. Tso. Human-robot interface using agents communicating in an xml-based markup language. In Procs. of the IEEE Int. Workshop on Robot and Human Interactive Communication, 2000. Osaka, Japan, September 27-29.
[14] A. Maslow. Motivation and Personality. New York Harper, 1954.
[15] John Peterson, Paul Hudak, and Conal Elliott. Lambda in motion: Controlling robots with haskell. In First International Workshop on Practical Aspects of Declarative Languages (PADL), January 1999.
[16] G. N Saridis. On the Theory of Intelligent Machines: A Comprehensive Analysis. International Journal of Intelligent Control and Systems, 1(1):3–14, 1996.
[17] P.S. Schenker, T.L Huntsberger, P. Pirjanian, and G.T. McKee. Robotic autonomy for space:cooperative and reconfigurable mobile surface systems. In Procs. of the 6th Int. SYmp. on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS'01), 2001. Montreal, Canada, June 18-21.
[18] J. Scholtz. Theory and evaluation of human robot interactions. In Procs. of the 36th Hawaii Int. Conf. on System Sciences, HICSS'03, 2003. Hawaii, January 5-8, USA,.
[19] J. Sequeira and M. I. Ribeiro. Behaviour-Based Cooperation Between Two Mobile Manipulators. Journal of Integrated Computer-Aided Engineering, 7(3):193–215, 2000.
[20] J. Sequeira and M. I. Ribeiro. A geometric approach to single and multiple robot control. In Proceedings of the 7th IFAC Symposium on Robot Control, Syroco 2003, Wroclaw, Poland, September, 1-3 2003.
[21] J. Sequeira and M. I. Ribeiro. A geometric approach to single and multiple robot control. In Preprints of the 5th IFAC/EURON Symposium on Intelligent Autonomous Vehicles, IAV2004, 2004. Lisbon, Portugal, July 5-7.
[22] G. Smirnov. Introduction to the Theory of Differential Inclusions, volume 41 of Graduate Studies in Mathematics. American Mathematical Society, 2002.
[23] Georgia Tech. Real-time cooperative behavior for tactical mobile robot teams - subsystems specification. Technical Report A002, Georgia Tech College of Computing and Georgia Tech Research Institute, October 1984.
[24] J. C. Willems. Paradigms and Puzzles in the Theory of Dynamical Systems. IEEE Transactions on Automatic Control, 36(3), March 1991.