

Robot Team Control: A Geometric Approach*

João Sequeira

jseq@isr.ist.utl.pt

Instituto Superior Técnico/Institute for Systems and Robotics,
Av Rovisco Pais 1, 1049-001 Lisboa, Portugal

M. Isabel Ribeiro

mir@isr.ist.utl.pt

Abstract

This paper describes the control of robot teams in the framework of Hilbert spaces. The proposed paradigm develops in two levels: (i) single robot control supported on a monotonic and non-expansive projection map defined on the configuration space, and (ii) team control supported on a supervision scheme over a set of neighboring relationships among the team-mates.

Each robot monitors its own neighboring relationships for relevant changes, accounting both for the distance among the robots and their relative motion, and adapts its motion to the rest of the team using a finite state automaton supervisor.

Simulation results on teams of 2D holonomic and cart robots are presented.

1 Introduction

This paper presents an approach to the control of robot teams supported on basic tools from the geometry of Hilbert spaces. The paper is organized in three main parts: (i) single robot control in the framework of differential inclusions, (ii) team control in a hybrid systems framework (with the continuous state verifying a differential inclusion), and (iii) experimental results.

In the framework of dynamical systems the i -th robot in a team with n members, moving in a configuration space $Q_i = \{q_i\}$, is represented by a dynamic system $\dot{q}_i(t) = f_i(q_i(t), u_i(t))$, with initial condition $q_i(0) = q_{i_0}$, $i = 1, \dots, n$, t being the time and $u_i \in U_i$ the control vector¹. When considered isolated from the rest of the team, the synthesis of the u_i that makes the robot follow a reference path or move towards a reference configuration is a classical robot control problem that have been widely studied, [7, 10, 15, 17]. For robot teams, approaches ranging from control theory, [13, 18], to discrete event systems, [8], and artificial intelligence, [1, 4, 6, 12], have been considered.

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

¹Without losing generality, this model can simply be assumed as the differential kinematics.

A large variety of single robot missions neither requires exact path following nor that the robot reaches a specific configuration. Instead, the robot is required to move within some bounded region in the free configuration space and/or to reach a goal region or a specific configuration. This problem has been considered within the Viability Theory framework, [2], that computes the control such that some a priori defined set is a viability domain, i.e., a region containing at least one solution trajectory to the control problem.

Similar considerations apply to teams of robots operating either under tight or loose constraints on the distance among the team members. For example, to avoid that the distance among team members grows above a pre-specified limit, the team may be required to span a limited workspace region while moving towards the goal.

The problem considered in this paper requires that the i -th robot reaches a goal set, $K_i \subset Q_i$. This set can be a priori defined to account for the mission specifications and modified as a consequence of the data acquired by on-board sensors during the mission.

The team control paradigm proposed in this paper encompasses three levels: (i) definition of the goal sets K_i to which the local (to each robot) control strategy drives the robot i , and (ii) a negotiation procedure, handled by a finite state automaton (FSA) at each robot, to adapt the motion strategy to the requirements imposed by the team neighboring constraints.

Once defined the goal set K_i and in the absence of motion constraints, the set $\Delta_{K_i}(q_i) = \{k - q_i, \forall k \in K_i\}$ defines the motion directions that drive the i -th robot directly towards K_i . In general, given the constraints imposed by the robot dynamics, only a subset of $\Delta_{K_i}(q_i)$ can be used in the computation of the controls. This subset results from the intersection between the feasible velocities, $f_i(q_i, U_i) \equiv \{f_i(q_i, u_i), u_i \in U_i\}$, defined by the robot dynamics, and the desired velocities, $\Delta_{K_i}(q_i)$, defined from the mission goals. This matching between the velocity sets can be expressed as a control problem within the framework of differential inclusions by $\dot{q}_i \in \{f_i(q_i, U_i) \cap \Delta_{K_i}(q_i)\}$.

The existence of solutions to differential inclusions

requires regularity conditions, namely various types of continuity or semicontinuity, for the map $F_i : Q_i \rightsquigarrow \{f_i(q_i, U_i) \cap \Delta_{K_i}(q_i)\}$, [3]. Given the dependence of F_i on the environment (through Δ_{K_i}), these are seldom verified. Furthermore, the co-domain of F_i is often the empty set. To overcome these two problems, this paper proposes a two step procedure to solve the robot control problem in the framework of differential inclusions: (i) pre-processing the acquired data on the environment to ensure the adequate properties and (ii) synthesize a control such that, whenever the intersection between the velocity sets is empty, $f_i(q_i, U_i)$ and $\Delta_{K_i}(q_i)$ converge to each other, in a sense to be defined ahead.

The paper is organised as follows: Section 2 presents sufficient conditions for the convergence between two sets, used to solve the single robot control problem modelled by a differential inclusion. Sections 3 and 4 outline the supervision strategies used over the single robot control strategy defined in Section 2, for multiple robots. Section 5 presents simulation examples for 2D holonomic and cart robot teams. Section 6 concludes the paper and points directions for further work.

2 Motion in Hilbert spaces

The existence of a solution for the control problem formulated in Section 1 requires the convergence between two velocity sets. This convergence can be expressed by having a measure of the distance between the two sets converging to 0. Given that the configuration space of real mobile robots is a Hilbert space, the metric induced by an inner product can be used to define a distance between sets. In particular, the following definition for the distance between two sets A and B is used.

Definition 1 (Distance between subsets of a normed space)

Let A and B be two subsets of a normed space. The distance between A and B is defined as

$$d(A, B) = \min_{x_a \in A, x_b \in B} \|x_a - x_b\| \quad (1)$$

where $\|\cdot\|$ stands for a norm. Although commonly referred as a distance, [9], (1) is not a metric. The distance between a point $x_a \in A$ and a set B is defined by making A a singleton.

The following theorem defines a projection map of a point onto a set. See [3] for a demonstration.

Theorem 1 (Best Approximation Theorem) Let G be a closed convex subset of a Hilbert space, X . There is a map $\pi : X \mapsto G$ and a unique element $\pi_G(x)$ verifying

$$\forall x \in X, \quad \|x - \pi_G(x)\| = \min_{y \in G} \|x - y\|$$

The map $\pi : X \mapsto G$ projects X into the goal set G . It is called a *best approximation projector* (b.a.p.) whenever the following two conditions are verified

$$\begin{aligned} \|\pi_G(x) - \pi_G(y)\| &\leq \|x - y\| && \text{non-expansivity} \\ \langle \pi_G(x) - \pi_G(y), x - y \rangle &\geq 0 && \text{monotonicity} \end{aligned} \quad (2)$$

with $\langle \cdot, \cdot \rangle$ standing for an inner product and $\|\cdot\|$ for the corresponding induced norm.

The generation of a trajectory converging towards a convex set can be derived using some basic tools of geometry. The following theorem states sufficient conditions for this convergence to occur. The demonstration follows from the use of a b.a.p., defined from Theorem 1 and the conditions (2), and of the Schwarz inequality (see [16] for details).

Theorem 2 (Point-to-set convergence) Let X be a Hilbert space, $G \subset X$ a convex set, $x \in X \setminus G$ and $\pi_G(x)$ a b.a.p. of a point x onto G . Furthermore, let x_n be a subsequence of a trajectory $x(t)$, outside G , such that

$$\exists T > 0 : \forall t_n > T, \langle x_{n+1} - x_n, \pi_G(x_{n+1}) - x_{n+1} \rangle > 0. \quad (3)$$

Then $\lim_{n \rightarrow \infty} d(x_n, G) = 0$.

Figure 1 illustrates an intuitive interpretation of Theorem 2 when X is the 2D euclidean space. The λ stands for the inner product in (3) and represents a rate at which the distance to G decreases by choosing the point x_{n+1} when the system is at x_n (assuming that controls driving the system between x_n and x_{n+1} exist).




Figure 1: Interpretation of Theorem 2 when X is a 2D Euclidean space

When $x_{n+1} \rightarrow x_n$ the condition in Theorem 2 can be formulated in continuous time, as

$$\langle \dot{x}, T_G^\perp(x) \rangle > 0 \quad (4)$$

where $T_G^\perp(x)$ represents the orthogonal contingent cone to G at x , i.e., the set of vectors orthogonal to the space tangent to the border of G at x .

It is worth to emphasize that no particular physical meaning was assigned to the X space and the goal set

G considered in this section. Theorem 2 represents a discrete time approach as x_n stands for a point sampled at time t_n . Inequality (4) expresses the continuous time equivalent convergence condition. For the purpose of this paper, the discrete time version is used with the identification $x_n \equiv \dot{q}_n$ whereas the continuous time version is used with the identification $x \equiv q$.

3 Control of single robots

Using the results in Section 2, the control of a single robot i amounts to compute the controls $U_{i_s} = \{u_i \in U_i : \lambda \geq 0\}$. Computing U_{i_s} may require the use of exhaustive search techniques to evaluate the projection map and/or to avoid solving inverse dynamics problems, e.g., computing u_{i_n} from $\dot{q}_{i_{n+1}}$.

In the absence of further criteria, a greedy strategy is used to choose a control vector from U_{i_s} with the chosen control vector being the one that yields the maximum λ_i , obtained from (3). This greedy approach may lead to local minima in $\lambda_i|_{U_{i_s}}$ for instance due to the loss of degrees of freedom in the control vector, i.e., one or more controls do not cause any progress in the mission even though the robot may be moving. To avoid being trapped in local minima, each robot has a FSA that modifies the motion strategy whenever $\lambda_i|_{U_{i_s}}$ approaches 0.

The convexity assumption on the goal set G_i (identified either with K_i or $\Delta_{K_i}(q_i)$ in the continuous or discrete cases, respectively) may require that a convex set is extracted out of the data acquired by the onboard sensors. Extracting a convex set out of an acquired raw G_i can be made using computational geometry techniques. For instance, assuming a simple polygonal region, polygon convexifying techniques, [5] or the convex hull, [14], can be used. An alternative technique is to decompose the raw G_i into Dirichlet regions, choosing one among them (see [11] for an introduction to the convex partitioning of simple polygons). Non polygonal convex regions strictly contained in G , with a shape that simplifies the computation of $\pi_{\Delta_{K_i}}(\dot{q}_{i_n})$ or $T_{K_i}^\perp(q_i)$, can also be used as a replacement for a raw goal set G .

4 Control of robot teams

The configuration of a team, given by $\{q_1, \dots, q_n\}$, determines a set of distance relationships among the team members (i.e., neighboring relationships). When these relationships are exact, the control of the team reduces to a standard set of n (one per robot) trajectory following control problems.

To increase the flexibility of the team in what concerns mission execution, the neighboring relationships may be relaxed, by allowing the distance between robots to vary within a specified range or according to some a priori defined mapping. The behavior of each

robot in the neighborhood of the teammates is determined by these relationships.

Neighboring relationships holding for the entire duration of a mission are topological invariants². For the purpose of this paper, a team of robots equipped with such a set of neighboring conditions is called a formation.

Clearly, arbitrary neighboring relationships may lead to non useful formations. For example, the orientation (in the world reference frame) of any robot in a team is a topological invariant under arbitrary translations. A team orientation, defined as a weighted sum of the individual orientations, is also invariant under arbitrary translations. However, translations defined as to make the robots collide are useless whereas translations preserving the distance between the robots may be useful.

Inter-robot distance preserving translations are a particular instance of a large class of formations characterized by having some region around each robot³ always in contact (e.g., intersecting or tangent) with the regions assigned to other robots. The neighboring relationships in this class are simply defined by specifying which regions of influence are in contact with which others. Among the practical cases of formations in this class are some soccer and military strategic movements.

The team control problem thus amounts to the computation of the motion of the team such that a set of a priori defined neighboring relationships is a topological invariant under the team motion. For a team of holonomic robots in the 2D plane, a practical topological invariant formations is the Voronoi diagram of the set of configurations $\{q_1, \dots, q_n\}$ that corresponds to the use of circular influence regions centered at each robot configuration. In this case, the above region defined around a robot can be identified with the Voronoi polygon the robot is in. If each of the polygons keeps in touch with the same set of neighbours along the whole mission the formation is preserved. Figure 2 illustrates two formation situations.


Given a set of neighboring conditions, the team control problem encompasses (i) the checking of the neighboring conditions and (ii) the synthesis of a FSA per robot to handle the corresponding events. The operation of this FSA corresponds to a negotiation procedure among the teammates.

5 Experimental results


This section presents simulation results on team control, illustrating the ideas outlined along the pa-

²A topological invariant is a mathematical object that maintains its qualities between any two initial and final configurations (e.g., connectedness) under the motion in the space of team trajectories for the assigned mission.

³The robot's region of influence.



Formation preserving the Voronoi polygons neighboring relationships



Formation preserving a short distance between teammates, but not the Voronoi polygons

Figure 2: Snapshot sequences illustrating two formations.

per. Two experiments are presented, both with teams of size 3. In both experiments, the mission assigned to the team is to move through a sequence of configurations. The mission ends as soon as one of the team members reaches its final goal region. In the first experiment the team is composed of 2D holonomic robots moving in the plane, each of which controlled according to the discrete strategy. In the second experiment the team is composed of cart robots, each of which controlled according to the continuous strategy. For the sake of simplicity, the robots are dimensionless in both experiments.

The goal sets K_i are balls of radius 0.1 centered in a priori defined locations (see Tables 1 and 2) whereas the sets of desired velocities are defined as $\Delta_{K_i}(q_i) = \{k - q_i, \forall k \in K_i\}$. To speed up the exhaustive search process in the computation of the λ function, the control space in both experiments is restricted to $\{(v, \omega)\} \equiv V \times \Omega$, with $V = \{-0.1, -0.05, 0, 0.05, 0.1\}$ and $\Omega = \{-0.1, -0.05, 0, 0.05, 0.1\}$. For the 2D holonomic robots, V and Ω stand, respectively, for the set of controls acting on the world reference frame (x and y axis). For the cart robots, V and Ω stand, respectively, for the set of linear and angular velocities in the robot reference frame. The b.a.p. map at instant t_n is given by $\pi_{\Delta_K}(\dot{q}_n) = -q_n + 0.1(q_n + q_n)/\|\dot{q}_n + q_n\|$ and, using Theorem 2, $\lambda = \langle \pi_{\Delta_K}(\dot{q}_n) - \dot{q}_{n+1}, \dot{q}_{n+1} - \dot{q}_n \rangle$. In the second experiment, $T_K(q)^\perp = -\alpha q/\|q\|$, with α an arbitrary positive constant, leading to $\lambda = \langle (\cos(\theta)v, \sin(\theta)v, \omega), \alpha(-q)'/\|q\| \rangle$.

Any control vector (v, ω) leading to $\lambda > 0$ is acceptable (assuming no other criteria than the convergence

conditions of Section 2). Control vectors leading to $\lambda = 0$ are acceptable in the discrete control strategy. The control vectors leading to $\lambda < 0$ are rejected.

The neighboring relationships between robots i and j considered in this paper are defined by the set of expressions

$$d(q_j, R_i) < 0.5 \quad (5)$$

$$d(q_j, R_i) > 0.1 \quad (6)$$

$$i = 1, \dots, n \quad j = \begin{cases} i+1 & \text{if } i < n \\ 1 & \text{if } i = n \end{cases}$$

where R_i stands for a convex influence region around the i -th robot. Condition (5) imposes that the robots do not move away from each other whereas condition (6) avoids any excessive approximation.

5.1 Team of 2D holonomic robots

The goal of the mission is to have the robots reaching the goal sets defined in Table 1 while preserving a reference formation. This reference formation is defined from the initial configuration of the robots and the neighboring conditions (5-6). These conditions allow each robot to move in a neighborhood of its teammates, with a preferred “follow the leader” motion strategy (even though the leader can change during the mission).

The simple kinematics considered allows the decoupling of the effects due to the negotiation procedure from those due to the single robot control strategy. For instance, in the absence of neighboring conditions, each of the robots proceeds straight to the respective goal set.

The FSA handling the negotiation is highly dependent on the particular assigned mission. In this experiment, the relevant states identify formation breaks and the reaching of goal sets. Any breaking of the reference formation must lead to a change in the motion strategy of each of the robots. The robots simply stop and wait for the teammates to resume the broken formation.

The mission is successfully completed. Figure 3a shows the trajectories of the robots, with the goal sets being represented by the circles. The slight bending of robot 1 trajectory near its second goal set indicates the end of a negotiation period and the beginning of a straight to goal motion. Figure 3b shows the time evolution of the λ function for each robot. The abrupt variations indicate the changes in the motion strategy defined by the negotiation FSA whereas the intervals with (positive) constant λ values represent the periods of time where controls are computed by specific, mission dependent, algorithms. The negative slopes following the peaks indicate the periods where the controls are computed by the procedure outlined in Section 3. Figure 3c shows the state evolution of the negotiation FSA. Holonomic robots are

Robot	Starting configurations	Goal sets center location		
		Goal set 1	Goal set 2	Goal set 3
1	(-0.8,-0.6)	(-0.8, 0.6)	(0.56, 0.59)	(0.5, -0.5)
2	(-0.7,-0.7)	(-0.7, 0.35)	(0.45, 0.42)	(0.5, -0.5)
3	(-0.6,-0.8)	(-0.6, 0.16)	(0.34, 0.27)	(0.5, -0.5)

Table 1: Mission specification for the 2D holonomic robots team

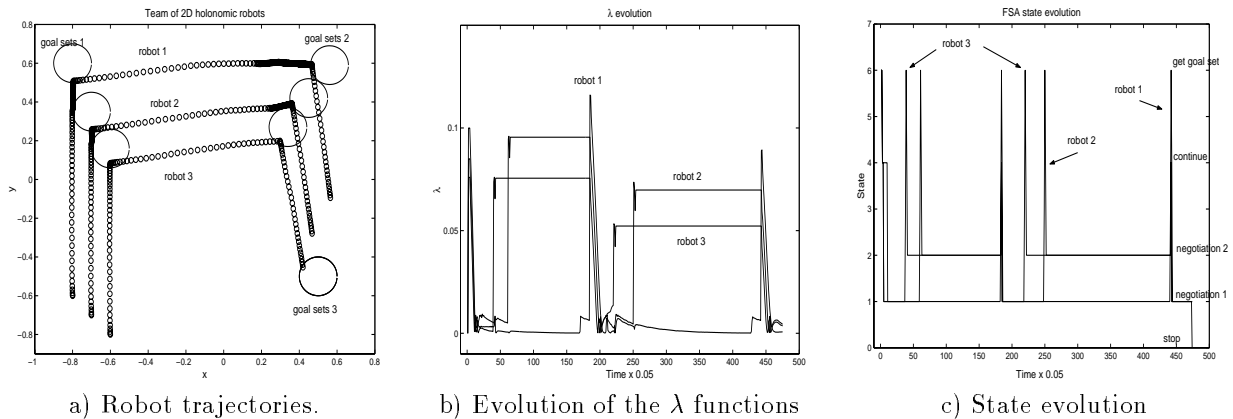


Figure 3: Mission results for the 2D holonomic robot team

guided straight to the goal by the single robot control algorithm. The neighboring conditions (5-6) avoid the robots either to get too close to each other or to break away. The team behavior depends on the individual control goal (i.e., moving towards a goal set) and of the chosen neighboring conditions. These have the effect of making relative positions of the robots approximately constant along the mission thus exhibiting a formation pattern.

5.2 Team of cart robots

The goal of the mission is to have the robots reaching the goal sets defined in Table 2 while preserving a reference formation. Only the neighboring relationship (5) is used in this mission. The resulting formation does not impose any order relationship among the teammates, meaning that each robot is free to move within a 0.5 distance of the teammates before a formation break occurs.

Figure 4 shows the mission results, with the goal sets shown as shaded regions in the 3D trajectories plot. The 3D trajectories show intense maneuvering as a result of the nonholonomy of the robots. This matches the behavior of the λ curves, with the spikes indicating the motion strategy changes at each robot, and of the negotiation FSA state, with multiple negotiations from all the teammates indicating the intense maneuvering. The mission was successfully completed, with robot 1 reaching the final goal set and

robots 2 and 3 shortly behind. The state evolution of the negotiation FSA exhibits a behavior more complex than in Experiment 1. This is mainly due to the fact that the intense maneuvering often required by cart robots tends to easily move the teammates apart.

6 Conclusions

This paper presented a geometric approach to the control of robots, supported on a projection map defined from basic properties of Hilbert spaces and sufficient conditions for the convergence of the robot configuration to a goal subset in the C -space.

The approach encompasses a two-level structure with: (i) an algorithm to choose controls verifying a sufficient condition for the mission execution under unconstrained motion, and (ii) a FSA to handle the interactions arising whenever the robots in a team compete to execute the mission.

Designing a formation with non constrained robots requires a set of neighboring relationships, a negotiation procedure among the teammates and the changes to the motion strategy of each robot. For non holonomic robots the analysis of a formation is far more complex. The region spanned by the team tends to be much larger than for holonomic robots due to the intense maneuvering that may occur along the mission.

Further work includes the study of stability issues for the proposed structure using results from hybrid systems theory and generalized Lyapunov functions.

Robot	Starting configurations	Goal sets center location		
		Goal set 1	Goal set 2	Goal set 3
1	(-0.8,-0.6,0)	(-0.31, 0.31, $\pi/2$)	(0.76, 0.47, 0)	(0.5, -0.5, $-\pi/2$)
2	(-0.7,-0.7,0)	(-0.56, 0.73, $\pi/2$)	(0.22, 0.48, 0)	(0.5, -0.5, $-\pi/2$)
3	(-0.6,-0.8,0)	(-0.79, 0.45, $\pi/2$)	(0.48, 0.76, 0)	(0.5, -0.5, $-\pi/2$)

Table 2: Mission specification for the cart robots team

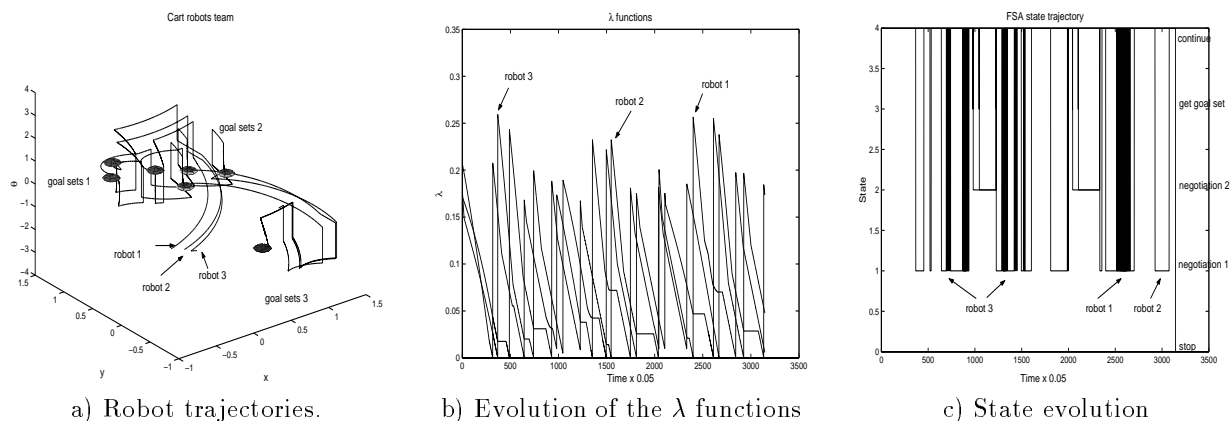


Figure 4: Mission results for the cart team

References

- [1] R. Arkin and T. Balch. Cooperative Multiagent Robotic Systems. In D. Kortenkamp, R. Bonasso, and R. Murphy, editors, *AI-Based Mobile Robots: Case Studies of Successful Robot Systems*. MIT Press, 1998.
- [2] J. P. Aubin. *Viability Theory*. Birkhäuser, 1991.
- [3] J.P. Aubin and A. Cellina. *Differential Inclusions*. Springer-Verlag, 1984.
- [4] A. Drogoull and A. Collinot. Autonomous Agents and Multi-Agent Systems. In *Applying and Agent-Oriented Methodology to the Design of Artificial Organizations: A Case Study in Robotic Soccer*. Kluwer Academic, 1998.
- [5] B. Grünbaum. How to convexify a polygon. *Geocombinatorics*, 5:24–30, 1995.
- [6] N. Jennings. Controlling Cooperative Problem Solving in Industrial Multi-Agent Systems Using Joint Intentions. *Artificial Intelligence*, (75):195–240, 1999.
- [7] O. Khatib, S. Quinlan, and D. Williams. Robot planning and control. *Robotics and Autonomous Systems*, (21):249–261, 1997.
- [8] J. Koščeká, H. Christensen, and R. Bajcsy. Experiments in behavior composition. *Robotics and Autonomous Systems*, 19:287–298, 1997.
- [9] E. Kreysig. *Functional Analysis*. Wiley, 1978.
- [10] R. Murray and S. Sastry. Nonholonomic Motion Planning: Steering using Sinusoids. *IEEE Transactions on Automatic Control*, 38(5), May 1993.
- [11] J. O’Rourke. *Computational Geometry in C*. Cambridge University Press, 2nd edition, 1998.
- [12] L. E. Parker. ALLIANCE: An Architecture for Fault Tolerant Multirobot Cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240, April 1998.
- [13] S. Ploen and F. Park. A Lie group formulation of the dynamics of cooperating robot systems. *Robotics and Autonomous Systems*, 21:279–287, 1997.
- [14] F. Preparata and M. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, 1985.
- [15] C. Samson. Path Following and Time-Varying Feedback Stabilization of a Wheeled Mobile Robot. In *Proceedings of the ICARCV’92*, Sept. 1992. Singapore.
- [16] J. Sequeira and M. I. Ribeiro. Geometric control of robot teams. Technical Report RT-602-03, Institute for Systems and Robotics / Instituto Superior Técnico, Lisboa, Portugal, 2003.
- [17] O. Sordalen. *Feedback Control of Nonholonomic Mobile Robots*. PhD thesis, The Norwegian Institute of Technology, 1993. Department of Engineering Cybernetics, N-7034 Trondheim, Norway.
- [18] Y. Yamamoto and X. Yun. Coordinating Locomotion and Manipulation of a Mobile Manipulator. In Zheng Yuan F, editor, *Recent Trends in Mobile Robots*. World Scientific, 1993.