

Collision Avoidance for Multiple Agent Systems

Dong Eui Chang
Department of Mechanical Engineering
University of California
Santa Barbara, CA, 93106-5070
dchang@engineering.ucsb.edu

Shawn C. Shadden, Jerrold E. Marsden, and
Reza Olfati-Saber
Department of Control and Dynamical Systems
California Institute of Technology 107-81
Pasadena, CA, 91125
{shawn,marsden,olfati}@cds.caltech.edu

Abstract— Techniques using gyroscopic forces and scalar potentials are used to create swarming behaviors for multiple agent systems. The methods result in collision avoidance between the agents as well as with obstacles.

I. INTRODUCTION

It is intriguing how swarms of insects or flocks of birds can travel in large, dense groups without colliding. Even in the presence of external obstacles these agents are capable of smoothly avoiding collisions. There is strong reason to believe that the rules, or protocols, each constituent of the group follows are quite basic, and yet the collective or global motion is quite remarkable. It comes as no surprise that understanding these protocols would be invaluable for engineering systems of autonomous agents, such as fleets of unmanned air or underwater vehicles or groups of exploratory robots.

The goal of this paper is to introduce a simple, decentralized control law that constituents of a group of vehicles can follow to accomplish some specified control objective while avoiding collision with one another and with unforeseen obstacles. In particular, we rely on the use of gyroscopic forces for collision avoidance, as described in Chang and Marsden [1]. Another paper that exploits gyroscopic forces Justh and Krishnaprasad [4], where the authors use gyroscopic forces to produce flocking behavior of multiple vehicles travelling at constant speed. Although we are not particularly impelled to reproduce flocking behavior, as in the famous work of Reynolds [8], the localized protocols that we developed largely to avoid collisions, do seem to create emergent, structured behavior when applied to large groups of vehicles.

Collision avoidance plays an important role in the context of managing multiple vehicles, especially in the context of air-traffic control (see the work of Tomlin and coworkers, [2], [3]). Many traditional control methods for collision avoidance rely on a potential-based approach, such as in the Navigation Function Method (NFM) of Rimon and Koditschek [9] or harmonic potential fields as in Masoud and Masoud [5]. The ideas presented in the paper are inspired by the NFM. The general idea of the NFM is to create a global potential field to accomplish some control objective, such as getting a vehicle to travel from its initial location to some target point while not colliding with any obstacles. To create this

global potential field, an attracting potential might be placed at the target point while repelling potentials are placed at the locations of obstacles to push an approaching vehicle away from the obstacle. Then the vehicle navigates using the gradient of the potential field as a force field.

The breakthrough of the NFM was that it could be used to show the existence of trajectories that avoid collision with any obstacles. However, this method has a few drawbacks: i) global information is needed regarding the location and shape of all the existing obstacles, ii) corresponding to any obstacle, there exists a neighborhood that can trap the vehicle for relatively long (or infinite) time, iii) the NFM is often computationally impractical, and iv) the original NFM only considers the case of a single vehicle.

Instead of relying on repelling potentials for obstacle avoidance, as in the NFM, the control law we present relies on gyroscopic forces. To motivate this, consider the situation shown in Fig. I, where three vehicles are initially equally spaced about a circle.

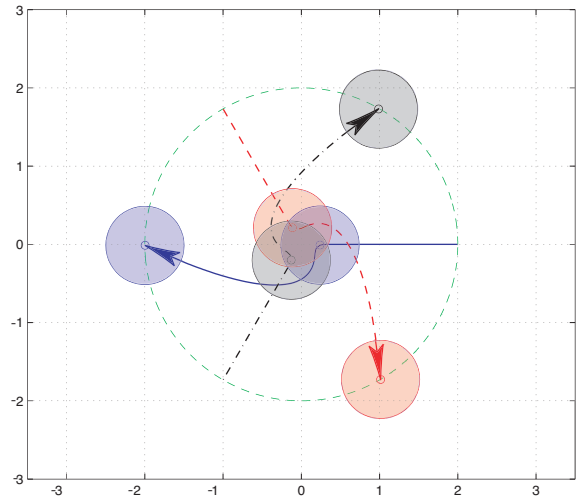


Fig. 1. By using gyroscopic forces, three vehicles can move across the circle without colliding. The vehicles are represented by the blue, black, and red dots. The disk about each dot represents the vehicle's detection shell.

As indicated in the figure, the objective of each vehicle is to simultaneously move to its antipodal point on the circle. If potential forces alone were used for collision avoidance,

the three vehicles would simply meet in the middle, push equally on each other, and become gridlocked. However, if gyroscopic forces are used, they can simply spin free of this sticky situation, as show in the figure.

Gyroscopic forces can be thought of as *steering* forces since they always act perpendicular to the direction of motion. Fortunately, gyroscopic forces can be used for obstacle avoidance without affecting the global potential function that was constructed to foster some control objective. Actually, gyroscopic forces do not even change the energy of the system, a well known and easily verified fact.

It turns out that gyroscopic forces alone have some difficulty preventing collision in large groups of vehicles. Therefore we also introduce a type of *braking* force that allows the vehicles to slow down when they are getting too close to one another or an obstacle. Intuitively, if the vehicle is moving too fast towards an obstacle it will not have enough time to turn to avoid the obstacle. Therefore the braking force is used to slow the vehicle such that it can turn in time to avoid the obstacle. As with the gyroscopic force, the braking force also does not change the global potential function.

The control law that we present is completely decentralized, therefore the ability of an agent to accomplish its control objective is not directly dependent on any other agent. Each vehicle has its own *detection shell*, within which it can sense the relative location of neighboring vehicles or obstacles. Since the control of each vehicle is localized, the computations can as well be localized, which is very important to ensure scalability of our control law to groups that contain a large number of agents. The methods that we present are equally applicable to either 2- or 3-dimensional motion.

In this paper we only consider “swarms” of vehicles. That is, we do not explicitly constrain the relative location of each vehicle. However, there are many applications where the relative positioning between vehicles becomes important, such as in NASA’s mission to construct a large interferometer from multiple telescopes (see <http://origins.jpl.nasa.gov> for more information). We refer to such groups of vehicles that are constrained to move in a particular pattern as *formations*. The authors plan to produce future work on collision avoidance in the context of formations.

The next section contains an overview of the development of our control law. We show figures that demonstrate the application of this control law to a collision avoidance scenario, known as the *split/rejoin* maneuver, coined by Olfati-Saber and Murray [6]. Since this paper only outlines preliminary results on collision avoidance and multi-vehicle maintenance, we will discuss some of our ongoing and future work. We conclude this paper by highlighting a few key points.

II. MAIN RESULTS

As previously mentioned, our goal is to develop a simple, decentralized control law that vehicles in a group can follow to maintain a collision free environment. Section II-A highlights the main points in the development of such a control law. To fully appreciate this section, one should refer to Chang and Marsden [1], where collision avoidance *except zero-velocity collision* is proved by means of gyroscopic, dissipative, and potential forces. A result related to the proof in [1] is given in the Appendix.

A. Collision Avoidance

Suppose we have a group of fully actuated vehicles obeying second-order translational dynamics (we do not consider attitude dynamics here). Since each vehicle will implement the same control law, we need only to develop the control for one vehicle.

Without loss of generality we can assume the vehicle has unit mass. We desire a feedback control law to (asymptotically) drive the vehicle to a target point \mathbf{q}_T without colliding with any obstacles or other vehicles. A *detection shell*, a ball of radius r_{det} , is given to the vehicle such that the vehicle can respond to any obstacle within this shell. For the purpose of designing the control law, let us refer to an obstacle of vehicle i as being *either* an external object that vehicle i is trying to avoid *or* a neighboring vehicle within vehicle i ’s detection shell. In addition, we assume all obstacles are convex.

The dynamics of the vehicle are given simply by $\ddot{\mathbf{q}} = u$, where $\mathbf{q} \in \mathbb{R}^3$. The control u consists of three parts as follows:

$$u = F_p + F_d + F_g \quad (1)$$

where F_p is a potential force which assigns to the vehicle a potential function with the minimum at the target \mathbf{q}_T ; F_d is a damping force; F_g is a gyroscopic force. The three forces, F_p , F_d , and F_g , are of the following form:

$$F_p = -\nabla V(\mathbf{q}), \quad F_d = -D(\mathbf{n})\dot{\mathbf{q}}, \quad F_g = S(\mathbf{n}, \dot{\mathbf{q}})\dot{\mathbf{q}}$$

where \mathbf{n} denotes the vector from the vehicle to its nearest obstacle (in particular, the nearest point on the obstacle), V is a (potential) function on \mathbb{R}^3 , the matrix D is symmetric and positive-definite, and the matrix S is skew-symmetric.

One suitable choice for the potential function is a simple quadratic $V(\mathbf{q}) = \frac{1}{2}\|\mathbf{q} - \mathbf{q}_T\|^2$. The matrix S in the gyroscopic force term $F_g = S(\mathbf{n}, \dot{\mathbf{q}})\dot{\mathbf{q}}$ is chosen to be an infinitesimal rotation about the vector $\mathbf{n} \times \dot{\mathbf{q}}$ when $\mathbf{n} \times \dot{\mathbf{q}} \neq 0$. When $\mathbf{n} \times \dot{\mathbf{q}} = 0$, a preferred rotational direction can be chosen. Since the matrix S is skew-symmetric,

$$\langle S(\mathbf{n}, \dot{\mathbf{q}})\dot{\mathbf{q}}, \dot{\mathbf{q}} \rangle = 0 \quad \Rightarrow \quad F_g \cdot \dot{\mathbf{q}} = 0.$$

Therefore F_g does not do any work.

The damping term, F_d , in (1) can be thought of as having two positive definite components:

$$F_d = -D(\mathbf{n})\dot{\mathbf{q}} = -(D_{dc} + D_b(\mathbf{n}))\dot{\mathbf{q}}$$

where D_{dc} is a constant matrix that represents any natural dissipation, such as drag, and $D_b(\mathbf{n})$ is an imposed *braking* term that varies with the relative distance between the vehicle and its nearest obstacle. It is this braking force, $D_b(\mathbf{n})\dot{\mathbf{q}}$, along with the gyroscopic force, that the vehicle uses for collision avoidance.

As previously mentioned, each vehicle is given a detection shell of radius r_{det} within which it can determine the relative location between itself and its nearest obstacle. Even if more than one obstacle is within the vehicle’s detection shell, *the vehicle only reacts to the nearest one*. If there is no obstacle within the vehicle’s detection shell then the gyroscopic and braking forces are zero. In addition, each vehicle does not react to obstacles “behind” them, even if the obstacle is within the detection shell. The magnitude of the gyroscopic and braking forces varies as a negative exponential of the distance between the vehicle and its nearest obstacle, for example:

$$D_b(\mathbf{n}) = C_1 \exp(-\|\mathbf{n}\|) - C_2$$

where C_1 and C_2 are positive constants chosen such that the force is bounded by some specified maximum and is zero when the obstacle is on the edge of the detection shell. Of course more elaborate reaction schemes can be devised, but the one described above is relatively very simple and requires only a minimal amount of communication between agents—two very desirable characteristics.

As with the NFM, we can prove, under some assumptions, that our control law can asymptotically drive a vehicle to its target point without the vehicle colliding with any obstacles (see [1] for details). The two main assumptions are that the kinetic energy of the vehicle is bounded and that only one obstacle is present in the detection shell. The second assumption is obviously the most restrictive and the one that we would like to relax, especially in the case for multiple vehicles. What is promising about using gyroscopic forces is that we have the ability to decouple, in a sense, the collision avoidance from the control objective. For instance, if we can prove the vehicle does not collide with any obstacles, then we can easily prove convergence of the vehicle to its target point using the energy $E(\mathbf{q}, \dot{\mathbf{q}}) = \frac{1}{2}\|\dot{\mathbf{q}}\|^2 + V(\mathbf{q})$ as a Lyapunov function since gyroscopic forces do not change the energy.

B. Demonstration of a Split/Rejoin Maneuver

Here we apply the control law developed in the previous section to a group of 25 vehicles. We provide only snap-shots of an interesting *split/rejoin* maneuver. To download and view this and related animations, please visit <http://www.cds.caltech.edu/~shawn/download.html>.




Fig. 2. Split/Rejoin maneuver for a group of 25 vehicles

Fig. 2 shows the evolution of 25 vehicles navigating around a large obstacle. For ease of visualization, the simulation is 2-dimensional, i.e. $\mathbf{q} \in \mathbb{R}^2$. At time $t = 0$, each agent starts off randomly positioned with an arbitrary velocity, Fig. 2a. The vehicles quickly coalesce into a distinct group before they encounter the obstacle, Figs. 2b and 2c. Figs. 2d and 2e show the group splitting and moving around the obstacle, after which, the vehicles rejoin into a coherent group again, as in Fig. 2f.

The “target point” for the vehicles in this simulation was actually chosen to have some dynamics. The vehicles are chasing a virtual vehicle that is moving in a straight line ahead of the group. If desired, once or if the vehicles near the target location, they should switch control laws to allow them to hover about the target location, as described in [4]. Although we do not specify or constrain the group to move in a particular pattern, we see an emergent structure present in, for instance, Figs. 2b and 2f, for which we can only explain with intuitive reasoning. However, we find the emergence of such behavior from seemingly nonrestrictive, local control laws to be quite interesting. We are investigating how local control protocols can be implemented to produce emergent patterns.

C. Future Directions

This paper is intended only to introduce some ideas on collision avoidance for multiple vehicles. There is clearly much more work to be done in the area of collision avoidance and in the area of multi-vehicle management in general.

One issue that we are currently working on is determining the best way to deal with *formations* or *pseudoformations*. We can loosely think of a formation as being a group of vehicles in which the relative locations of the vehicles is important (see Olfati-Saber and Murray [6] for a precise

definition). In [6] it is shown that for a group of n agents in a plane, at least $2n - 3$ inter-agent distance-based constraints are needed to construct a *rigid formation*. In addition, more constraints need to be specified to avoid the *folding phenomenon* in rigid structures, as in Olfati-Saber, Dunbar, and Murray [7]. Applications where the use of formations as rigid and unfoldable graphs are important often arise when a group of vehicles forms a sensor network. For the purpose of flocking, a looser definition of a formation as a *flexible graph* seems to be more appropriate.

In the interest of managing large groups of vehicles, we are also interested in studying the role of hierarchal infrastructures and heterogenous groups. In the simulations shown in this paper, all vehicles were modelled identically and each vehicle applied the same control law. We would like to explore circumstances where vehicles in a group have different capabilities, objectives, preferences, or roles.

III. SUMMARY

The area of multi-agent control has many unsettled problems. It was our goal in this paper to stress an alternative approach for collision avoidance, namely the use of gyroscopic and braking forces. We were particularly attracted to these types of forces for several reasons, including: they do not affect the potential function; they seem more natural than forces derived from inter-agent potentials; and they are typically easier to implement on real-world systems since most vehicles have the ability to steer and brake but are not usually able to be pushed in arbitrary directions. Since the reaction protocols are based only on local information, they easily extend to groups with an arbitrary number of agents.

IV. ACKNOWLEDGMENTS

This research was supported by ONR/AOSN-II contract N00014-02-1-0826 through Princeton University.

APPENDIX

Here we suggest an algorithm for collision avoidance *using only gyroscopic forces*. Consider a (point-mass) vehicle with the following dynamics:

$$\ddot{\mathbf{q}} = \begin{bmatrix} 0 & -u \\ u & 0 \end{bmatrix} \dot{\mathbf{q}}$$

with $\mathbf{q} = (x, y)^T \in \mathbb{R}^2$. Let $V_o \neq 0$ be the initial speed of the vehicle. Since the gyroscopic force does not change the speed, we have $\|\dot{\mathbf{q}}(t)\| = V_o$ for all $t \geq 0$. Assume the gyroscopic force is bounded by the positive constant ω_o ; in particular we allow $u \in \{0, \pm\omega_o\}$. Hence when $u = 0$, the motion is linear, and when $u = \omega_o$ (resp. $-\omega_o$), the vehicle moves counter-clockwise (resp. clockwise) on a circle of radius V_o/ω_o . We assume that the vehicle has a detection shell so that it can detect the relative location of any obstacles within a circular neighborhood of radius R_d .

We now present a simple algorithm which allows a vehicle to approach a target point, avoiding any obstacles on the journey.

a) Reorientation.: First we give the reorientation algorithm, which orients the vehicle so that it points to the target. Let \mathbf{n}_t denote the vector from the vehicle to the target point, and θ_t denote the signed angle from $\dot{\mathbf{q}}$ to \mathbf{n}_t . We choose u following the rule:

$$u = \begin{cases} \omega_o & \text{if } \theta_t > 0, \\ 0 & \text{if } \theta_t = 0, \\ -\omega_o & \text{otherwise.} \end{cases} \quad (2)$$

Hence the vehicle aligns its velocity toward the target unless it senses an obstacle in its detection shell.

b) Collision Avoidance.: We now describe an algorithm for avoiding collision between the vehicle and an obstacle. Assume that the obstacle is a convex body, B . Let $B_e = \{\mathbf{q} \in \mathbb{R}^2 \mid d(\mathbf{q}, B) \leq R_d\}$, where $d(\cdot, \cdot)$ is the Euclidean distance function.

Suppose the vehicle has detected an obstacle, i.e. $\mathbf{q} \in B_e$. Let \mathbf{n}_o denote the vector from the vehicle to the obstacle and θ_o be the signed angle from \mathbf{n}_o to $\dot{\mathbf{q}}$. The collision avoidance algorithm is as follows:

$$u = \begin{cases} \omega_o & \text{if } \theta_o > 0 \\ -\omega_o & \text{otherwise} \end{cases} \quad (3)$$

Then, in a finite time, the vehicle will leave the region, B_e , assuming $R_d > V_o/\omega_o$ (that the vehicle's turning radius is less than the width of the detection shell). This inequality is somewhat conservative. One can prove lower bounds under appropriate restrictions. As an example, for circular obstacles one can show that if

$$V_o \leq \frac{\omega_o(R_d^2 + 2r_o R_d)}{2R_d}$$

where r_o is the radius of the obstacle, no collision will occur.

c) Combination.: We would like to combine the reorientation algorithm (2) and the collision avoidance algorithm (3) so that the vehicle can ultimately hit the target, while avoiding collisions with any obstacles. We assume that obstacles are at least more than $2(R_d + V_o/\omega_o)$ away from one another. The algorithm can be succinctly specified as follows:

ALGORITHM I:

1 Use (2)

2 IF obstacle is in the detection shell

THEN use (3)

ELSE goto 1

Assuming obstacles are located more than $2(R_d + V_o/\omega_o)$ apart, no limit cycles will occur due to the introduction of multiple obstacles with ALGORITHM I. Although we do not yet have an analytic convergence proof for this algorithm, a

geometric proof exists, but is omitted for brevity. A simple depiction of a trajectory specified by ALGORITHM I is given in Fig. 3.




Fig. 3. Collision avoidance with a bang-bang gyroscopic force.

II. REFERENCES

- [1] D.E. Chang and J.E. Marsden, "Gyroscopic Forces and Collision Avoidance", *Proc. of Conference in Honor of A. J. Krener's 60th Birthday*, 2002.
- [2] I. Hwang and C. Tomlin, "Multiple aircraft conflict resolution under finite information horizon", *Proceedings of the AACC American Control Conference*, Anchorage, May 2002, and "Protocol-based conflict resolution for air traffic control", *Stanford University Report SUDAAR-762*, July, 2002.
- [3] J. Kosecká, C. Tomlin, G. Pappas, and S. Sastry, "Generation of conflict resolution maneuvers for air traffic management", *IROS*, 1997; see also *IEEE Trans. Automat. Control* 43, 1998, pp 509-521.
- [4] E.W. Justh and P.S. Krishnaprasad, "A Simple Control Law for UAV Formation Flying", *ISR Technical Report 2002-38*, 2002.
- [5] A.A. Masoud and S.A. Masoud, "Motion Planning in the Presence of Directional and Obstacle Avoidance Constraints Using Nonlinear, Anisotropic, Harmonic Potential Fields", *Proc. IEEE ICRA*, San Francisco, CA, 2000.
- [6] R. Olfati-Saber and R.M. Murray, "Graph Rigidity and Distributed Formation Stabilization of Multi-Vehicle Systems", *Proc. of the 41st Conference on Decision and Control*, 2002.
- [7] R. Olfati-Saber, W.B. Dunbar, and R.M. Murray, "Cooperative Control of Multi-Agent Systems using Cost Graphs and Optimization", *Proc. of the American Control Conference*, 2003.
- [8] C.W. Reynolds, "Flocks, Herds, and Schools: A Distributed Behavioral Model, in *Computer Graphics*", *SIGGRAPH '87 Conference Proceedings*, 21(4), 1987, pp 25-34.
- [9] E. Rimon and D.E. Koditschek, "Exact robot navigation using artificial potential functions", *IEEE Trans. on Robotics and Automation*, 8(5), 1992, pp 501-518.