5-23-2008

Pareto optimal multi-robot coordination with acceleration constraints

Jae Bum Jung
University of Illinois at Urbana-Champaign

Robert Ghrist
University of Pennsylvania, ghrist@seas.upenn.edu

Pareto optimal multi-robot coordination with acceleration constraints

Abstract
We consider a collection of robots sharing a common environment, each robot constrained to move on a roadmap in its configuration space. To program optimal collision-free motions requires a choice of the appropriate notion of optimality. We work in the case where each robot wishes to travel to a goal while optimizing elapsed time and consider vector-valued (Pareto) optima. Earlier work demonstrated a finite number of Pareto-optimal classes of motion plans when the robots are subjected to velocity bounds but no acceleration bounds. This paper demonstrates that when velocity and acceleration are bounded, the finiteness result still holds for certain systems, e.g., two robots; however, in the general case, the acceleration bounds can lead to continua of Pareto optima. We give examples and explain the result in terms of the geometry of phase space.

Keywords
Pareto optimisation, mobile robots, multi-robot systems, path planning, Pareto optimal multirobot coordination, acceleration constraints, multirobot motion planning

2008 IEEE International Conference on
Robotics and Automation
Pasadena, CA, USA, May 19-23, 2008

Pareto optimal multi-robot coordination with acceleration constraints

Jae Bum Jung
Department of Mathematics
University of Illinois
Urbana IL, USA
jaejung@math.uiuc.edu

Robert Ghrist
Department of Mathematics
Coordinated Science Laboratory
University of Illinois
Urbana IL, USA
ghrist@math.uiuc.edu

***Abstract*— We consider a collection of robots sharing a common environment, each robot constrained to move on a roadmap in its configuration space. To program optimal collision-free motions requires a choice of the appropriate notion of optimality. We work in the case where each robot wishes to travel to a goal while optimizing elapsed time and consider vector-valued (Pareto) optima.**

Earlier work demonstrated a finite number of Pareto-optimal classes of motion plans when the robots are subjected to velocity bounds but no acceleration bounds. This paper demonstrates that when velocity and acceleration are bounded, the finiteness result still holds for certain systems, e.g., two robots; however, in the general case, the acceleration bounds can lead to continua of Pareto optima. We give examples and explain the result in terms of the geometry of phase space.

I. Introduction

This note considers the well-studied problem of multi-robot motion planning. The specific focus is on (vector-valued, or Pareto-) optimal coordination of agents, each of which has elapsed time-to-goal as its cost function. Such optimal coordination problems introduce challenges beyond those of simple obstacle avoidance and motion planning in single robot settings; robot-robot interactions must likewise be controlled.

A. Motivation

Our perspective is to emphasize vector-valued optimization, preserving all cost function data. This notion of **Pareto optimality** [20], [23] is standard in mathematical economics to model individual consumers striving to optimize distinct economic goals. It avoids data loss that comes with scalarization: *e.g.*, minimizing average time, or total elapsed time. Such scalarizations are both common and commonly appropriate in robotics ([13], [18], [26]), yet there is a loss associated with this scalar reduction. In the context of, say, a dynamic manufacturing or warehousing scenario, the priorities associated to individual agents may change from day-to-day, resulting in ever-changing optimization problems.

This note treats the global optimization problem for multiple robot coordination without scalarizing the vector-valued cost function. This centers on the notion of *Pareto optimality* [20], [23], a concept which is widely used in mathematical economics to model individual consumers striving to optimize distinct economic goals. The classification of Pareto optima automatically yields the set of all optima for all (monotone) scalarizations of the cost functions: see [11]. In addition, it provides a (hopefully small) template of optimal coordinations which can be used for on-line adaptation to changing needs and cost functions in day-to-day factory operations.

Given the desire to filter the space of all possible coordination schemes to a small set of best cases independent of biases on the robots, we are certainly most interested in the cases where this collection of optima is *finite*, and the existence of such is the focal point of this note.

B. History

Multi-robot coordination is of course a special case of general motion planning for multiple robots, for which a long history of work exists. Centralized approaches typically construct paths in a configuration space derived from the Cartesian product of the configuration spaces of the individual robots (e.g., [2], [3], [24]). More decoupled approaches may generate independent robot paths and then resolve illegal interactions between the robots (e.g., [5], [9], [19]).

The approach in [9] prioritizes the robots, and defines a sequence of planning problems for which each problem involves moving one robot while those with higher priority are considered as predictable, moving obstacles. This involves the construction of two-dimensional path-time space [14] over which the velocity of the robot is tuned to avoid collisions with the moving obstacles. In [1], [4], [7], [21], [19], [25] robot paths are independently determined, and a coordination diagram is used to plan a collision-free trajectory along the paths. The approaches in [1], [21] additionally consider dynamics. In [16], [27], an independent roadmap is computed for each robot, and coordination occurs on the Cartesian product of the roadmap path domains. In [17], an approximate Dijkstra-like algorithm to find Pareto optimal solutions was given. The suitability of one approach over the other is usually determined by the trade-off between computational complexity associated with a given problem, and the amount of completeness that is lost. In some applications, such as the coordination of AGVs, the roadmap might represent all allowable mobility for each robot.

There are very few results which give a rigorous classification of Pareto optima [8], [11], [10]. The paper [11] gives a finiteness result for Pareto optima with respect to elapsed

This work supported by DARPA # HR0011-07-1-0002.

978-1-4244-1647-9/08/$25.00 ©2008 IEEE.

time in the setting of AGVs restricted to roadmaps. However, these results assumed only a bound on velocity, not on acceleration. Acceleration constraints have been considered in several important works for scalar optimization. In [6], an exact algorithm for kinodynamic planning in the 2-d case was given: many of the ideas here are crucial in our analysis. In [22], the first known polynomial-time approximation algorithm for curvature-constrained shortest-path problems in higher dimensions was given.

C. Contributions

This note extends the finiteness results of [11] in the context of acceleration bounds. Section II reviews the classification of Pareto-optima in the bounded velocity case [11]. In §III, we argue that an appropriate first step for implementing acceleration bounds is to impose an upper bound without a lower bound: acceleration is more limited than deceleration. We show that in this context, the finiteness results for Pareto-optimal paths persists for two robots. In §IV we give a canonical example of a system with three robots for which the acceleration upper bound forces an infinite collection (a continuum in fact) of inequivalent Pareto-optimal path classes. We conclude with a geometric explanation for this change in behavior from 2-d to 3-d: it is regulated by the discrete curvature of the system's phase space.

II. Review: unbounded acceleration case

This section contains basic definitions and a review of the finiteness theorem for Pareto optimal robot coordinations from [11].

A. Coordination spaces

Recall that each robot travels on a **roadmap**, represented as a 1-d subspace or graph, Γ_i, $i = 1, \ldots, N$. The coordination space is the product of these roadmaps with all illegal or collision sets removed.

Definition 2.1: A **roadmap coordination space** of graphs $\{\Gamma_i\}_1^N$ is any space of the form

$$\mathcal{X} = \left(\prod_{i=1}^{N} \Gamma_i\right) - \mathcal{O}, \qquad (1)$$

where $\mathcal{O}$ denotes an (open) obstacle set.

For simplicity, one may focus on the case where each factor Γ_i is a single edge. This is the case where a robot translates along a track from an initial to final location. All illustrations in this note follow this convention, though the results hold for the general setting. All coordination spaces are assumed to be sufficiently "tame" (see [11] for details).

Most coordination spaces arising in robotics have an obstacle set $\mathcal{O}$ which is cylindrical, in the following sense:

Definition 2.2: A coordination space $\mathcal{X}$ is said to be **cylindrical** if $\mathcal{O}$ is of the form

$$\mathcal{O} = \bigcup_{i<j} \left\{ (x_k)_1^N \in \prod_{k=1}^{N} \Gamma_i \,:\, (x_i, x_j) \in \Delta_{i,j} \right\}, \qquad (2)$$

for some (open) sets $\Delta_{i,j} \subset \Gamma_i \times \Gamma_j$ where $1 \le i < j \le N$.

That is to say, a cylindrical coordination space is one for which illegal states are determined by pairwise configurations. If two robots have collided, it makes no difference what the positions or configurations of the remaining robot are — this state still counts as an illegal "collision" state.

B. Pareto-optimality

A coordination of N robots is a path γ in the roadmap coordination space $\mathcal{X}$. Throughout this note, each robot will use its elapsed time τ_i as a cost function with which optimality is measured. However, since there are N robots, there is a **cost vector** $\tau = (\tau_1, \ldots, \tau_N)$ that is a function of the coordination γ.

A path $\gamma : [0, T] \to \mathcal{X}$ is **pareto optimal** if and only if $\tau(\gamma)$ is minimal with respect to the partial order on vectors:

$$\tau(\gamma) \le \tau(\gamma') \Leftrightarrow \tau_i(\gamma_i) \le \tau_i(\gamma_i') \quad \forall\, i = 1 \ldots N. \qquad (3)$$

Two paths γ and γ' are Pareto equivalent if and only if they are homotopic through locally Pareto optimal paths which are equal in the partial order; *i.e.*, $\tau(\gamma) = \tau(\gamma')$, see Fig. 1.

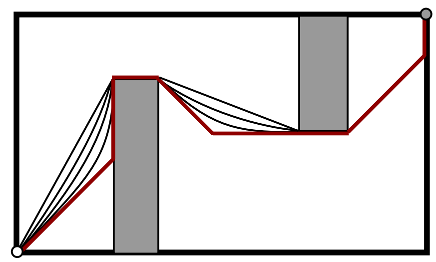

Fig. 1. A collection of Pareto-optimal paths weaving through obstacles forms a single equivalence class. Each roadmap is equipped with own metric and parametrized.

Theorem 2.3 ([11]): On a simply-connected cylindrical coordination space, there is a unique Pareto-optimal (bounded velocity) path class between fixed endpoints. On a general cylindrical coordination space, the number of globally Pareto-optimal path classes is finite.

The key to Theorem 2.3 is the construction of a canonical Pareto-optimal path: such *left-greedy* paths are reviewed in the subsequent section.

III. Initial bounded acceleration

The finiteness result of §II relies crucially on the lack of bound on acceleration, since the left-greedy paths used as canonical path classes *always* involve sudden starts and stops. Of course, acceleration bounds are critical in any reasonable robot system, and the question becomes to what extent these constraints effect Pareto-optima. If we add the acceleration bound then the class of admissible paths becomes much more complex, making optimization more challenging.

A. Left-greedy paths

Assume that $\mathcal{X}$ is a simply-connected (connected and 'hole-free') cylindrical coordination space and $p, q \in \mathcal{X}$ are fixed endpoints for a coordination. Let γ be a path from p to q. For any point $y = (y_k)_1^N \in \mathcal{X}$ on the path γ, consider

Algorithm 1 $(\boldsymbol{x}, V) = \text{IBALeftGreedy}(\mathcal{X}, \gamma)$

Require: γ is a collision-free coordination in $\mathcal{X}$.

1: Using γ label each obstacle, determine critical events, and compute crossing sequence.
2: Start at the initial point x_0, and set $i = 1$.
3: Apply maximum or minimum acceleration until the integral curve meet critical events from Step 1 and store the velocity profile $V(t)$.
4: Let $x_i - x_{i-1} + \int_{T_{i-1}}^{T_i} V(s)ds$ where $[T_{i-1}, T_i]$ is the time interval from x_{i-1} to x_i.
5: If x_t is not the goal point, then increment i and go to Step 3. Otherwise, terminate and report $\boldsymbol{x} = (x_0, \ldots, x_t)$ and $V(s)$.

the N distinct **hyperplanes** at y: $\mathcal{H}_k(y)$ is defined to be the connected component of $\{x \in \mathcal{X} : x_k = y_k\}$ containing y.

The following definition comes from [10].

Definition 3.1: A path in $\mathcal{X}$ from p to q is **left-greedy** if it crosses all hyperplanes separating p and q as quickly as possible. More specifically, for any $y \in \gamma$ and all $k = 1...N$, the (forward) tangent vector to the path, $\dot{\gamma}(y)$, satisfies the following:

(1) If $\mathcal{H}_k(y)$ separates p from q in $\mathcal{X}$, then the k^{th} component of $\dot{\gamma}(y)$ is nonzero and is positive/negative so as to point from p to q.
(2) If $q \in \mathcal{H}_k(y)$, then the k^{th} component of $\dot{\gamma}(y)$ is zero;
(3) All components of $\dot{\gamma}(y)$ are maximized with respect to the speed constraints of 1 and the obstacle constraints.

In [11], [10] it was shown that left-greedy paths form a canonical representative of the unique Pareto-optimal path class between fixed endpoints on a simply-connected cylindrical coordination space. For non-simply-connected coordination spaces, one can restrict attention to **homotopy classes** of paths (the universal cover is simply-connected).

Unfortunately, left-greedy paths are not of bounded acceleration: condition (3) certainly violates the bounded acceleration constraint. Therefore we need to put more conditions to (3) say that all components of $\dot{\gamma}(y)$ are maximized with respect to the speed constraints, the acceleration constraints, and the obstacle constraints. We will define an equivalent 'smoothed' version.

In a manufacturing/automation situation, there is a sharp distinction between acceleration and deceleration phases of motion: it is easier to stop than to go. For example, a standard factory AGV weighs more than a human and, barring the presence of an uncommonly large and energy-draining engine, fast accelerations are difficult. However, quick decelerations (especially when the terminal velocity is reasonably low) are much easier, being obtainable at the expense of heat generation and wear (friction on brakes) or mechanical means (bumpers).

For the remainder of this note we therefore assume an initial acceleration bound: there is a fixed upper bound on positive acceleration, but not for deceleration.

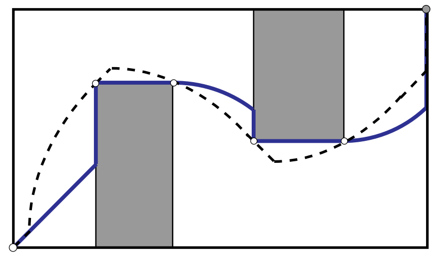

Fig. 2. An Initial Bounded Acceleration left-greedy path (solid) and a Pareto-optimal path (dashed). Empty circles are OC (off-contact) points.

1) IBA left-greedy paths: Algorithm 1 computes an IBA (Initial Bounded Acceleration) left-greedy path. It start from initial point with zero velocity vector. We can decide the maximum velocity and acceleration we need at the current point by the information from Step 1, and keep moving forward or backward until it meet the next critical event. The sign of acceleration depends on hyperplanes that separate the current point and the goal point. One continues until the path reaches the goal point. During this step, one stores the velocity profile V and the critical points x_t. The bounded-velocity left-greedy path is a piecewise-linear path for which robots maintain constant velocity on each segment. The IBA left-greedy path is a concatenation of segments for which robots maintain almost-constant acceleration at each segment.

A BV left-greedy path is a canonical representative of a Pareto optimal path [11], [10]. An IBA left-greedy path is not Pareto optimal in general. The difference is that whenever some robot restarts from a rest position it consumes more time than the robot which follows the Pareto optimal path, *i.e.*, the velocity at this instant is submaximal. We call this point the Off-Contact (OC) point (see Fig. 2). We generate a Pareto optimal path by modifying velocities at OC points. The next algorithm computes the path that has maximum or minimum velocities at OC points.

B. *Critical paths for two robots*

For the remainder of this section we assume that $\mathcal{X}$ is a two-dimensional coordination space.

1) Algorithms: Roughly speaking, OC points are points that can have better velocities on an IBA left-greedy path. We can detect OC points by checking the velocity profile. The velocity profile $V(t)$ is a vector valued function $V(t) = (V_1(t), V_2(t))$. At the OC point, one of the functions V_i must start increasing or decreasing its velocity from 0. Therefore we need to insert the line that detects OC points in Algorithm 1. Suppose $x_t = (x(1), x(2))$ is an OC point.

We say i is an **unsaturated direction** for x_t if $V_i(t)$ is changed from zero to nonzero — the other direction is called the **saturated direction** [6]. Let X be the critical point if the next OC point has the different saturated direction.

Let X_1 be a first critical point and t_1 be a time from the initial point to X_1. The time t_1 only depends on the saturated direction. Replace $V(t_1)$ by $\tilde{V}(t_1)$ such that (1) there exists

Algorithm 2 CRITICALPATH

1: Let y_0 be the initial point and $j = 1$.
2: Start algorithm 1.
3: Stop the algorithm when x_i is an OC point. Let $y_j = x_i$.
4: If the saturated direction is not changed from y_{j-1} then let $x_0 = x_i$, increment j and goto Step 2 otherwise compute the maximum velocity at y_{j-1}, store the critical point and velocity profile corresponding to the maximum velocity.
5: Let $x_0 = y_{j-1}$, increment j, and goto Step 2.

a admissible path from the initial point to X_1 with a velocity $\tilde{V}(t_1)$; (2) the velocity of unsaturated direction is maximized or minimized among all paths satisfying (1). Thus we have a new velocity $\tilde{V}$ at X_1. Once it hits the one of the goal position we are now in the one-dimensional space therefore the next goal time only depends on the velocity at this goal position which already is maximized. We also consider the goal positions are critical points. Finally we get path which is a concatenation of optimal segments connecting critical points. We call this new path $\tilde{\gamma}$ the critical path. The critical path is not unique in general but they share the same critical points and goal times. So we can form a equivalence class of critical paths.

2) Pareto optimality: We now prove that the critical path γ is actually a canonical Pareto-optimal path. Let $\mathcal{H}_s(a)$ be the hyperplane such that $\mathcal{H}_s(a) = \{x \in \mathcal{X} : x_1 = a_s$, where a_s is a saturated direction$\}$ and $\mathcal{H}_u(a)$ be the hyperplane such that $\mathcal{H}_u(a) = \{x \in \mathcal{X} : x_1 = a_u$, where a_u is an unsaturated direction$\}$ [6].

Lemma 3.2: Suppose a_0 is a critical point on γ and a_1 is the next critical point which is not the goal position such that $\gamma(t_0) = a_0$ and $\gamma(t_1) = a_1$. Suppose b is any point on $\mathcal{H}_s(a_1)$ and there exists an optimal path β which is homotopic (fixing endpoints) to γ such that $\beta(t_0) = a_0$, $\dot{\beta}(t_0) = \dot{a_0}$, and $\beta(t_1) = b$.

Then (1) $b_s = a_{1s}$ and $\dot{b_s} = \dot{a}_{1s}$. (2) If $\dot{b_u} > \dot{a_u}$ then there exists $t_2 > t_1$ such that $\gamma(t_2) = \beta(t_2)$ and $\dot{\gamma}(t_2) = \dot{\beta}(t_2)$ otherwise γ and β never meet at the same time.

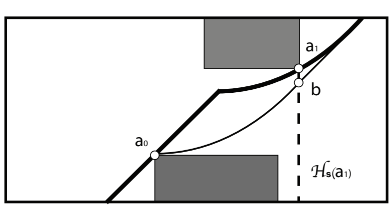

Fig. 3. The thick line is the path γ and the thin line is the path β.

Proof: Let a_1 be (a_{1s}, a_{1u}). (1) Trivially $b_s = a_{1s}$ since $b \in \mathcal{H}_s(a_1)$. Also t_1 totally depends on a_{1s} and $\dot{a}_{1s}$ if $\dot{b_s} \neq \dot{a}_{1s}$ then $b \notin \mathcal{H}_s(a_1)$.

(2) Assume $\dot{a_u} < v_{max}$ otherwise it is trivial. The maximal (or minimum) velocity at a_u depends on Δx_a which is the distance that robot traveled in the unsaturated direction and $\Delta t = t_1 - t_0$. Since $\dot{b_u}$ is greater than $\dot{a}_{1u}$ if we continuously change the velocity $\dot{a_u}$ to $\dot{b_u}$ then either Δt or Δx must be changed by maximality. Since Δt is fixed by the assumption we can only change Δx. The difference $|a_u - b_u|$ is proportional to the difference of a distance: see Fig. 3. ■

The following is the first principal result of this note: in dimension two, Pareto-optima obey a finiteness result even in the case of (initial) bounded acceleration.

Theorem 3.3: Suppose $\mathcal{X}$ is a 2-dimensional roadmap coordination space. Suppose γ is a critical path and β is another path which is homotopic to γ in $\mathcal{X}$. Then the cost vector $\tau(\gamma)$ of γ is less than or equal to the cost vector $\tau(\beta)$ of β: *i.e.*, any critical path γ is a Pareto optimal path.

Proof:

Suppose β is any path with the same endpoints as γ. Without a loss of generality, we may assume β is the concatenation of optimal segments (since, if not, we can replace with optimal segments at no increase in cost vector).

Suppose $\mathcal{A} = \{a_1, ..., a_M\}$ is a set of critical points of γ. Let $\mathcal{H} = \{\mathcal{H}_1, ..., \mathcal{H}_M\}$ be a set of hyperplanes where $\mathcal{H}_i = \mathcal{H}_s(a_i)$, $i = 1, ..., M$.

$\mathcal{H}_i$ separates $\mathcal{X}$ into two connected pieces; one that contains an initial point and another containing a goal point. Therefore, β must pass all $\mathcal{H}_i$, $i = 1, ..., M$ at least once. Also because $\mathcal{H}_i$ and $\mathcal{H}_{i+1}$, $i = 1, ..., M-1$, are defined by different saturated directions, it is safe to assume β follows the same order that γ do when it passes $\mathcal{H}$. See Fig. 4

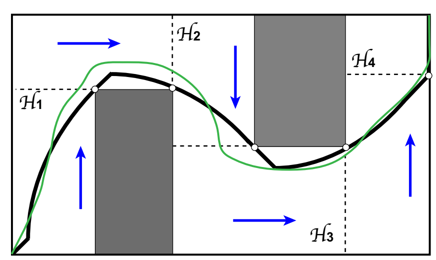

Fig. 4. γ (thick line) and β (thin line) go through hyperplanes. Arrows indicate saturated directions.

It is clear that $\tau(\gamma) = \tau(\beta)$ if β hit every critical points of γ. Suppose then that β does not pass through all critical points of γ and $\gamma(t_0) = a_j$ is the last critical point before two paths break down. Let $\gamma(t_a) = a_{j+1} = (a_1, a_2)$ be the next critical point of γ and let $\mathcal{H}_j$ and $\mathcal{H}_{i+1}$ be corresponding hyperplanes in $\mathcal{H}$ for a_j and a_{j+1}, respectively.

We focus on comparing the cost of segments of γ and β connecting $\mathcal{H}_j$ and $\mathcal{H}_{j+1}$. When a time T is a cost of γ for the segment, there are only three cases of β; the cost of β is less than T, equal to T, and greater than T.

The first case is fail due to the fact a time T totally depends on a saturated direction of a_{j+1}. Decreasing T implies changing velocity or a position of a_j and it contradicts that β also hits a_j. The second case can be explained by Lemma3.2. In this case, β either catchs up γ without losing any cost or arrives at next critical point of γ with a bigger total time.

For the last case, suppose $b = (b_s, b_u)$ is the point on β which lies on $\mathcal{H}_{j+1}$, where b_u is the component correspond to the unsaturated direction of a_{j+1}. The increment of time never decrease the cost of a coordinate b_s. Thus we only need to look at b_u. Assume $\dot{a_u} < 1$ where a_u is a component in saturated direction of a_{j+1} unless the cost is always greater than γ. Let v_m be the maximal velocity. Then we can compute time t_m which we need to achieve v_m from $\dot{a_u}$. By optimality assume $t_u \leq t_m$. Suppose β has the best possible velocity at b corresponding to the additional time t_u. But γ also takes t_u for the unsaturated direction to reach the same velocity β has at b and saturated direction must be changed at the next critical point. Thus β can not arrive in concurrence with γ at the next critical point of γ.

So we only need to consider the case they never meet again before the goal position. But, trivially, the time expended on the saturated direction is increased and the best that the unsaturated direction can do is arriving at the goal position at the same time as γ, *i.e.*, after t_u the unsaturated direction reaches its goal position. Therefore the cost vector of β is greater than the cost vector of γ. ∎

Corollary 3.4: Any other Pareto-optimal path in $\mathcal{X}$ which is homotopic to a critical path is Pareto equivalent to a critical path.

Proof: Suppose γ is a critical path and $\tilde{\gamma}$ is a Pareto-optimal path which is homotopic (fixing the endpoints) to γ. Since $\tau(\gamma) \leq \tau(\tilde{\gamma})$ by Theorem 3.3, $\tau(\gamma)$ must be equal to $\tau(\tilde{\gamma})$ and the only path that shares the same cost vector is a path that passes through every critical point of γ or case (2) in Lemma 3.2. ∎

IV. N-DEGREES OF FREEDOM

A cylindrical coordinate space $\mathcal{X}$ can be described by a set of 2-d projections. Therefore when we only have a bounded velocity constraint we can easily extrapolate algorithms from the 2-d case. Unfortunately, the argument fails in the case of bounded acceleration.

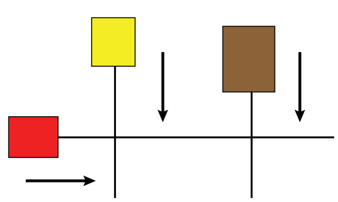

Fig. 5. There are two connected obstacle sets in the 3-d coordination space which are in different projections and doesn't intersect each other.

Theorem 4.1: Suppose $\mathcal{X}$ is a cylindrical coordinate space and admissible paths are paths that have bounded (or initial bounded) acceleration. Then there is no finite bound on the number of globally Pareto optimal classes in general.

Proof: We show this by example. Suppose there are 3 robots in the workspace, as in Fig. 5. Then $\mathcal{X}$ is a 3-dimensional cylindrical coordinate space. Suppose β is an IBA left-greedy path. Let $x = (x_0, y_0, z_0)$ be the unique critical point of β with the velocity $\dot{x} = (1, 0, 1)$. The 0-velocity in the second component is due to the obstacle in (x, y)-plane: see Fig. 6. Suppose the maximal admissible velocity of x is $(1, v_0, 1)$ which means that there is a path γ such that $\gamma(t_0) = x$ and $\dot{\gamma}(t_0) = (1, v_0, 1)$. Clearly $v_0 < 1$ because of the obstacle in the (z, y)-plane. If we slow down along the z-direction, then v_0 can be increased: see Fig. 7. So we can define a map $\Pi(t, s)$ such that $\Pi(t_0, 0) = \gamma(t_0)$, $\Pi(t_0, 1) = (x_0, y_0, z_1), z_1 < z_0$, and $\dot{\Pi}(t_0, 1) = (1, 1, 1)$. Then $\Pi(t, s_0)$ gives a one-parameter family of Pareto optimal paths, the cost vectors of which are $(c, k - h, l + h)$ where c, k, l are constant and $h > 0$. This continuum of paths is therefore pairwise inequivalent. ∎

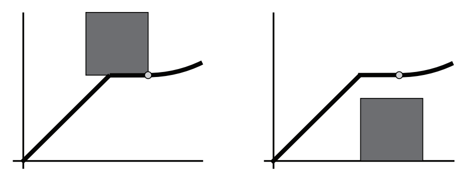

Fig. 6. IBA left-greedy path on the (x, y) plane [left] and (x, z) plane [right].

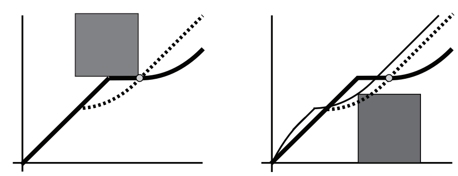

Fig. 7. The dotted line on the right indicates a collision path; the thin line is modified to avoid the collision by modifying the robot on the z-axis.

The difference between the bounded and unbounded acceleration cases is how obstacle sets are defined in the phase space. In the case of unbounded acceleration, a configuration space and a phase space have the same kind of obstacle set; both are cylindrical. But if we add the bounded acceleration constraint the obstacle set in a phase space becomes much more complicated. Obstacles are comprised of: 1) cylindrical obstacles (from collisions); 2) the region of inevitable collisions that depends on the speed; and 3) time-limited unreachable sets. Unlike types 1) and 2) type 3) depends on the path end points (initial and goal points) crucially.

3) Unbounded acceleration case: We first demonstrate that unbounded acceleration leads to a cylindrical coordination phase space.

Suppose $X = \prod_{i=1}^{N} \Gamma_i$ and $\mathcal{O}$ is a obstacle set which is cylindrical. Let $\mathcal{PX}$ denote the phase space $X \times \dot{X}$ where $\dot{X} = \{(v_1, ..., v_N) \,|\, \|v_i\|_\infty \leq 1\}$. Let us check the possible obstacles in $\mathcal{PX}$. Suppose $x \times \dot{x} \in \mathcal{PX}$ and $x \in \mathcal{O}$ then $x \times \dot{x}$ clearly is in the obstacle of $\mathcal{PX}$ which is cylindrical. Also if $x \times \dot{x} \in \mathcal{O}$ then x must be in $\mathcal{O}$. Therefore a configuration space $\mathcal{PX} - \mathcal{PO}$ is a cylindrical coordinate space.

4) Bounded acceleration case: The above result fails in the bounded acceleration case.

Now we must consider the region of inevitable collision, denoted by X_{ric}. X_{ric} is a set containing points in obstacle sets and also points which can not avoid the future colision[15]. For example, suppose $\mathcal{X}$ is 2-dimensional space. For $(x, y, 1, 1) \in X_{ric}$, there exist x', y', v_x, and v_y such that $(x', y, 1, 1), (x, y', 1, 1), (x, y, v_x, 1)$, and $(x, y, 1, v_y)$ are not in X_{ric}: see Figure 8. Therefore there exist noncylindrical obstacle sets.[1]

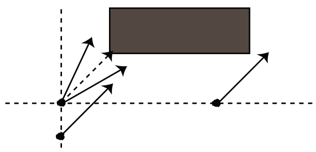

Fig. 8. Only the dotted arrow is in X_{ric}

V. Conclusion

The bounded velocity assumption on paths used in earlier works [11], [10] yields a very clean mathematical theory for classifying Pareto-optimal paths for multi-robot coordination. This note gives the first results for the bounded acceleration case. We restrict to the case of initial bounded acceleration to respect the physical differences between acceleration and deceleration in robotics. The two principal results are as follows.

(1) In the case of two robots, initial bounded acceleration does not alter the finiteness results for Pareto-optimal path classes.

(2) In the case of three or more robots which are sufficiently 'entangled' — which come close enough to each other to have obstacles in the coordination space which are not well-separated — the acceleration bounds force multi-parameter continua of distinct Pareto-optimal path classes.

In addition, we have observed that the cylindrical constraints on the appropriate coordination space (noted in [11] to be of fundamental importance to the finiteness results) are satisfied for IBA phase spaces in the 2-d case and are violated in higher dimensions. This lends credence to the proposition that cylindricity (and nonpositive curvature associated with it) is a fundamental reason for the (surprising) finiteness results.

Future work consists of determining bounds on obstacle separation to ensure a finiteness result in general multi-robot coordination problems. In addition, one can address the problem of terminal acceleration bounds to determine what, if any, effect these have on Pareto-optima. More broadly, the general problem of computing the topological type of the space of Pareto-optimal paths in robot coordination problems where the agents are unconstrained is both open and challenging.

References

[1] S. Akella and S. Hutchinson. Coordinating the motions of multiple robots with specified trajectories. In *Proc. IEEE Int. Conf. on Robot. and Autom.*, pages 624–631, 2002.

[2] M. D. Ardema and J. M. Skowronski. Dynamic game applied to coordination control of two arm robotic system. In R. P. Hämäläinen and H. K. Ehtamo, editors, *Differential Games - Developments in Modelling and Computation*, pages 118–130. Springer-Verlag, Berlin, 1991.

[3] J. Barraquand and J.-C. Latombe. Robot motion planning: A distributed representation approach. *Int. J. Robot. Res.*, 10(6):628–649, December 1991.

[4] Z. Bien and J. Lee. A minimum-time trajectory planning method for two robots. *IEEE Trans. Robot. & Autom.*, 8(3):414–418, June 1992.

[5] S. J. Buckley. Fast motion planning for multiple moving robots. In *IEEE Int. Conf. Robot. & Autom.*, pages 322–326, 1989.

[6] J. Canny, A. Rege and J. Reif, An exact algorithm for kinodynamic planning in the plane, in *Proc. Symp. on Comp. Geom.* 271–280, 1990.

[7] C. Chang, M. J. Chung, and B. H. Lee. Collision avoidance of two robot manipulators by minimum delay time. *IEEE Trans. Syst., Man, Cybern.*, 24(3), 517–522, 1994.

[8] H. Chitsaz, J. M. O'Kane, and S. M. LaValle. Pareto optimal coordination of two translating polygonal robots on an acyclic roadmap. In *Proc. IEEE International Conference on Robotics and Automation*, 2004.

[9] M. Erdmann and T. Lozano-Perez. On multiple moving objects. In *IEEE Int. Conf. Robot. & Autom.*, pages 1419–1424, 1986.

[10] R. Ghrist, J. M. O'Kane, and S. M. Lavalle, Computing Pareto optimal coordinations on roadmaps, *Intl. J. Robotics Research*, 12(11), 997–1010, 2005.

[11] R. Ghrist and S. M. Lavalle, Nonpositive curvature and Pareto-optimal coordination of robots, *SIAM J. Control & Opt.*, 45(5), 1697–1713, 2006.

[12] A. Hatcher., Algebraic Topology. *Cambridge University Press*, 2001

[13] H. Hu, M. Brady, and P. Probert. Coping with uncertainty in control and planning for a mobile robot. In *IEEE/RSJ Int. Workshop on Intelligent Robots and Systems*, pages 1025–1030, Osaka, Japan, November 1991.

[14] K. Kant and S. W. Zucker. Toward efficient trajectory planning: the path-velocity decomposition. *Intl. J. Robotics Research* 5(3), 72–89, 1986.

[15] S. M. Lavalle, Planning Algorithms, *Cambridge Univ. Press*, 2006.

[16] S. M. LaValle and S. A. Hutchinson. Path selection and coordination of multiple robots via Nash equilibria. In *Proc. 1994 IEEE Int'l Conf. Robot. & and Autom.*, pages 1847–1852, May 1994.

[17] S. M. LaValle and S. A. Hutchinson. Optimal motion planning for multiple robots having independent goals. *IEEE Trans. on Robotics and Automation*, 14(6):912–925, December 1998.

[18] J. Miura and Y. Shirai. Planning of vision and motion for a mobile robot using a probabilistic model of uncertainty. In *IEEE/RSJ Int. Workshop on Intelligent Robots and Systems*, pages 403–408, Osaka, Japan, May 1991.

[19] P. A. O'Donnell and T. Lozano-Pérez. Deadlock-free and collision-free coordination of two robot manipulators. In *IEEE Int. Conf. Robot. & Autom.*, pages 484–489, 1989.

[20] L. E. Parker. Cooperative motion control for multi-target observation. In *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, pages 1591–1598, 1998.

[21] J. Peng and S. Akella. Coordinating multiple robots with kinodynamic constraints along specified paths. In *Algorithmic Foundations of Robotics V, STAR 7*, Springer-Verlag, 221–237, 2004.

[22] J. Reif and H. Wang, Nonuniform discretization for kinodynamic motion planning and its applications, *SIAM J. Comput.* 30(1), 161–190, 2000.

[23] Y. Sawaragi, H. Nakayama, and T. Tanino. *Theory of Multiobjective Optimization*. Academic Press, New York, NY, 1985.

[24] J. T. Schwartz and M. Sharir. On the piano movers' problem: III. Coordinating the motion of several independent bodies. *Int. J. Robot. Res.*, 2(3):97–140, 1983.

[25] T. Simeon, S. Leroy, and J.-P. Laumond. Path coordination for multiple mobile robots: a resolution complete algorithm. *IEEE Trans. Robot. & Autom.*, 18(1), February 2002.

[26] S.-H. Suh and K. G. Shin. A variational dynamic programming approach to robot-path planning with a distance-safety criterion. *IEEE Trans. Robot. & Autom.*, 4(3):334–349, June 1988.

[27] P. Svestka and M. H. Overmars. Coordinated motion planning for multiple car-like robots using probabilistic roadmaps. In *IEEE Int. Conf. Robot. & Autom.*, 1631–1636, 1995.

[1]Indeed, from the perspective of [11], there is a discrete form of positive curvature associated to these coordination phase spaces. Such positive curvature is the culprit in destroying the finiteness results for Pareto-optima.