

NONPOSITIVE CURVATURE AND PARETO-OPTIMAL COORDINATION OF ROBOTS*

ROBERT GHRIST[†] AND STEVEN M. LAVALLE[‡]

Abstract. Given a collection of robots sharing a common environment, assume that each possesses a graph (a 1-d complex also known as a ROADMAP) approximating its configuration space and, furthermore, that each robot wishes to travel to a goal while optimizing elapsed time. We consider vector-valued (or PARETO) optima for collision-free coordination on the product of these roadmaps with collision-type obstacles. Such optima are by no means unique: in fact, continua of Pareto optimal coordinations are possible.

We prove a finite bound on the number of optimal coordinations in the physically relevant case where all obstacles are CYLINDRICAL (*i.e.*, defined by pairwise collisions). The proofs rely crucially on perspectives from geometric group theory and CAT(0) GEOMETRY. In particular, the finiteness bound depends on the fact that the associated COORDINATION SPACE is devoid of positive curvature.

We also demonstrate that the finiteness bounds holds for systems with moving obstacles following known trajectories.

1. Introduction.

1.1. Motivation. In numerous settings, the coordination of multiple robots remains a basic and challenging research issue. Autonomous guided vehicles (AGVs) are used in a wide variety of industrial settings for problems such as material handling, palletizing, paper roll handling, assembly, conveying, and people moving. Typically, AGVs reliably traverse a fixed roadmap of paths in a complicated factory environment. Although the paths avoid collisions with obstacles in the workspace, the efforts of numerous AGVs may have to be coordinated in a way that avoids collisions between AGVs while at the same time maximizing productivity.

If we wish to coordinate the motions of N robots in a common environment, what is an appropriate notion of optimality? Minimizing the average time robots take to reach their goal? Minimizing the time that the last robot takes? Such approaches are common (e.g., [19, 25, 34]) and may be appropriate in some cases; however, it is important to recognize that scalarization of a vector of N criteria occurs in this process. Each robot has its own cost function, e.g., elapsed time. These N criteria are then converted — often in an arbitrary manner — into a single criterion to be optimized.

In this paper, we investigate the optimization problem for multiple robot coordination without scalarizing the vector-valued cost function. This centers on the notion of *Pareto optimality* [29, 31], a concept which is widely used in mathematical economics to model individual consumers striving to optimize distinct economic goals. This brand of optimization is richer in the sense that no choice of scalarization is involved. In particular:

1. Any optimum of *any* scalarization of the vector-valued cost function (which is monotone increasing in the components) is in fact one of the Pareto optima: see

*Supported by NSF PECASE DMS-0337713 [RG] and DARPA # HR0011-05-1-0008 [RG,SL].

[†] Department of Mathematics and Coordinate Science Laboratory, University of Illinois, Urbana, IL 61801, USA

[‡] Department of Computer Science, University of Illinois, Urbana, IL 61801, USA

Lemma 2.6.

2. For settings in which priorities change over time, knowledge of the Pareto optima allow for easy adaptation. This is particularly relevant in automation settings where priority changes arise from, say, changing needs for parts delivery or output.
3. By filtering out all of the motion plans that are not worth considering and presenting the user with a small set of the best alternatives, additional criteria, such as priority or the amount of sacrifice one robot makes, can be applied to automatically select a particular motion plan.

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*. A finiteness criterion is the principal result of this paper.

1.2. Contributions. In §2 we will phrase the coordination problem as an optimal path-planning problem on a ROADMAP COORDINATION SPACE \mathcal{X} : a product of finite metric graphs $\{\Gamma_i\}_1^N$ with obstacle sets removed. The optimization criteria are distinct for each robot. Thus, the appropriate category of optimal paths are the so-called PARETO OPTIMAL solutions. Roughly speaking, a path between fixed endpoints is LOCALLY PARETO OPTIMAL if and only if deforming it to any path which decreases one of the goal times increases some other goal time. A path is GLOBALLY PARETO OPTIMAL if and only if any other path which decreases one goal time increases some other. Paths are PARETO EQUIVALENT if and only if they are connected by a 1-parameter family of paths which preserve all goal times; this equivalence generates Pareto optimal CLASSES. (See §2.2 for precise definitions).

Depending on the coordination space, inequivalent Pareto optimal classes may be isolated or may appear in uncountable continuous families. Examples appear in §3. Most desirable for classification and cataloging purposes are those coordination spaces which admit a finite number of globally Pareto optimal classes.

The principal optimization result of this paper is that for coordination spaces which possess ‘cylindrical’ obstacles — those defined by pairwise collisions between robots — there is a complete classification of locally Pareto optimal classes and a finite bound on the global Pareto optima.

Main Theorem: Let \mathcal{X} be a cylindrical coordination space:

1. Fixing endpoints, the locally Pareto optimal path classes are in bijective correspondence with homotopy classes of paths.
2. There is a finite bound on the number of globally Pareto optimal path classes.

The first statement implies that the set of Pareto optimal paths within a fixed homotopy classes of paths is connected in this space of paths. The second statement means that only a finite number of these local optima are in fact globally Pareto optimal. Proofs appear in §5.

We furthermore extend the finiteness bound to a class of systems involving moving obstacles: see §7. Remarkably, our techniques easily adapt to problems for which obstacles of varying shape, speed, and trajectory can interfere with the robot motion.

The proof of the Main Theorem uses in an essential manner certain techniques from geometric group theory: namely, the geometry of CAT(0) SPACES (see, *e.g.*, [8]) and certain

NORMAL FORMS for groups acting on CAT(0) cubical complexes (as in [26]). These tools are applicable to the proof of the Main Theorem thanks to:

Main Lemma: Let \mathcal{X} be a cylindrical coordination space. Then \mathcal{X} is locally CAT(0) .

This theorem is a broad generalization of the result by Abrams [1] that configuration spaces of graphs are locally CAT(0) in the product metric.

Though it falls outside the scope of this paper, the CAT(0) techniques we develop here are directly applicable to the problem of *computing* normal forms of Pareto optimal paths. See [17] for an exact algorithm for computing Pareto optimal paths in a cylindrical coordination space.

1.3. History. The problem of coordinating robots along fixed roadmaps can be considered as a special case of general motion planning for multiple robots. Previous approaches to multiple-robot motion planning are often categorized as *centralized* or *decoupled*. A centralized approach typically constructs a path in a composite configuration space, which is derived from the Cartesian product of the configuration spaces of the individual robots (e.g., [5, 6, 32]). A decoupled approach typically generates paths for each robot independently, and then considers the interactions between the robots (e.g., [9, 15, 27]). The results of this paper are for decoupled coordination.

The approach in [15] 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 [20] over which the velocity of the robot is tuned to avoid collisions with the moving obstacles.

In [3, 7, 11, 30, 27, 33] robot paths are independently determined, and a coordination diagram is used to plan a collision-free trajectory along the paths. The approaches in [3, 30] additionally consider dynamics. In [22, 35], an independent roadmap is computed for each robot, and coordination occurs on the Cartesian product of the roadmap path domains. 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 automated guided vehicles, the roadmap might represent all allowable mobility for each robot.

2. Definitions.

2.1. Coordination spaces. This paper concerns the coordination of N robots, each having a roadmap Γ_i (a graph within the configuration space of the i^{th} robot) pre-computed independently of the other robots. All graphs are assumed finite and possessed of an intrinsic metric (inherited from the configuration space in which it resides).

DEFINITION 2.1. A ROADMAP COORDINATION SPACE of graphs $\{\Gamma_i\}_1^N$ is any space of the form

$$\mathcal{X} = \left(\prod_{i=1}^N \Gamma_i \right) - \mathcal{O}, \tag{2.1}$$

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

ASSUMPTION 2.2. *For the remainder of the paper, all coordination spaces are assumed to have obstacle sets which are sufficiently nice. Namely, the boundary of the obstacle set, $\partial\mathcal{O}$, is piecewise real-analytic and COLLARED: there exists some embedding of $\partial\mathcal{O} \times [0, 1] \hookrightarrow \mathcal{X}$ which takes $\partial\mathcal{O} \times \{0\} \rightarrow \partial\mathcal{O}$ via the identity map.*

This assumption rules out obstacle sets \mathcal{O} for which the system “locks up” in a singular configuration. The piecewise smoothness condition for $\partial\mathcal{O}$ is imposed to avoid discussions of *chattering* and non-rectifiable paths.

EXAMPLE 2.3. A special class of a roadmap coordination space arises when all of the graphs Γ_i are identical and the obstacle set is an open neighborhood of the collision set $\{x_i = x_j : \text{for some } x_i \in \Gamma_i; x_j \in \Gamma_j; i \neq j\}$. In this case, one can consider the workspace to be the graph itself, and the roadmap coordination space is precisely the configuration space of N labeled points on the graph.

In practice, a large class of coordination spaces arise by assigning as illegal states \mathcal{O} those configurations for which there is a collision between robots in the common workspace. For such an system, the obstacle set \mathcal{O} has a very particular “cylindrical” structure since all collisions are determined by pairwise data.

DEFINITION 2.4. *A coordination space \mathcal{X} is said to be CYLINDRICAL if \mathcal{O} is of the form*

$$\mathcal{O} = \bigcup_{i < j} \left\{ (x_k)_1^N \in \prod_{k=1}^N \Gamma_k : (x_i, x_j) \in \Delta_{i,j} \right\}, \quad (2.2)$$

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

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.

Configuration spaces of graphs are examples of cylindrical coordination spaces, since the obstacle set \mathcal{O} is the union of (neighborhoods of) sets of the form $\{x_i = x_j : i \neq j\}$. Non-cylindrical coordination spaces arise when the legality of a potential collision between two robots is dependent upon the configuration of a third: e.g., hallways in the workspace through which two but not three robots can squeeze.

2.2. Pareto optimality. Pareto optimization refers to optimization of vector-valued functions [28]. In the context of robotics applications, Pareto optimization arises when distinct robots possess distinct goals and cost functions for evaluating performance. Each robot wishes to optimize its cost function independently of the others.

Mathematically, this is characterized as follows. Given a parameterized path $\gamma : [0, T] \rightarrow \mathcal{X}$ in a coordination space, each robot executes the projected path $\gamma_i = \text{PROJ}_i \circ \gamma$, where PROJ_i denotes projection onto the i^{th} factor. Given cost functions $\{\tau_i\}$, $i = 1 \dots N$, the cost vector for γ is the vector $\tau(\gamma) = (\tau_i(\gamma_i))_i^N$. The case in which τ_i measures *elapsed time* from start to goal is an important and characteristic example, though more general cost functions are allowed.

ASSUMPTION 2.5. *We restrict attention to parameterized paths which are ADMISSIBLE, meaning that velocity vectors have all components bounded above by one. There are no restrictions on acceleration: only velocity is bounded. It will be assumed for simplicity that*

the cost functions τ_i agree with elapsed time. To change the velocity bound or the cost function, one can perform a simple re-scaling of the geometry of the coordination space.¹

A path $\gamma : [0, T] \rightarrow \mathcal{X}$ is PARETO OPTIMAL iff $\tau(\gamma)$ is minimal with respect to the partial order on vectors:

$$\tau(\gamma) \leq \tau(\gamma') \Leftrightarrow \tau_i(\gamma_i) \leq \tau_i(\gamma'_i) \quad \forall i = 1 \dots N. \quad (2.3)$$

The local Pareto optima comprise the set of all optima for all monotone scalarizations (such as, e.g., average time-to-goal and non-linear generalizations thereof):

LEMMA 2.6. *For any scalarization $f : \mathbb{R}^n \rightarrow \mathbb{R}$ with $\partial f / \partial x_i > 0$, all minima of $f \circ \tau$ are global Pareto optima.*

Proof: Given any minimizer for $f \circ \tau$ which is not globally Pareto optimal, there exists a path which decreases some τ_i without increasing any of the others; by monotonicity, this decreases the f -value of the path. Contradiction. \diamond




FIG. 2.1. [left] The unique locally and globally Pareto optimal path between corners of a rectangle with regards to elapsed time; [right] an envelope of Pareto optimal paths weaving through obstacles forms a single equivalence class.

In some cases, locally Pareto optimal paths are unique. In the rectangular coordination space of Fig. 2.1[left], there is a unique locally (and globally) Pareto optimal path. The diagonal of the rectangle is not locally Pareto optimal unless the rectangle is a square.

Locally Pareto optimal paths are usually not unique. Two paths γ and γ' are PARETO EQUIVALENT iff they are homotopic through locally Pareto optimal paths which are equal in the partial order; i.e., $\tau(\gamma) = \tau(\gamma')$. Fig. 2.1[right] illustrates a single locally Pareto optimal class with many representatives. We show in §3 that certain roadmap coordination spaces admit a continuum of locally/global Pareto optimal *classes*.

2.3. A little topology. Experts can safely bypass this brief primer of the topological concepts used throughout the paper. See, e.g., [18] for a comprehensive treatment.

Given a topological space X , a PATH is a (continuous) map $\alpha : [0, 1] \rightarrow X$. Two paths α and α' sharing identical endpoints are said to be HOMOTOPIC if there exists a continuous 1-parameter family of paths $\{\alpha_t\}_0^1$ satisfying (1) $\alpha_0 = \alpha$; (2) $\alpha_1 = \alpha'$; and (3) α_t has the same endpoints for all t .

¹The assumption about potentially unbounded acceleration is not so easily skirted. The results of this paper are unknown in that case.

Topologists keep track of different homotopy classes of paths by computing a more general object: the fundamental group. Given a fixed basepoint $x_0 \in X$, the FUNDAMENTAL GROUP, $\pi_1(X, x_0)$, is defined to be the set of homotopy equivalence classes of paths in X which have endpoints at x_0 . The group operation is concatenation of paths: $\alpha * \beta$ means “do α then do β .” This is well-defined (since all endpoints are x_0) and yields a group operation with identity element the CONTRACTIBLE loops in X . For path-connected spaces, $\pi_1(X, x_0)$ is independent of the basepoint x_0 (up to isomorphism of groups) and one may unambiguously denote the fundamental group $\pi_1(X)$. A SIMPLY CONNECTED space X is one which is connected and has $\pi_1(X)$ a trivial group.

Any space X which has a reasonably nice local structure (e.g., a cell complex) and a non-trivial fundamental group can be “unwrapped” to a larger simply connected space which is locally indistinguishable from X . The UNIVERSAL COVER of X , \tilde{X} , is defined to be a simply connected space along with a projection map $p : \tilde{X} \rightarrow X$ which is locally a homeomorphism. For example, the universal cover of the circle S^1 is \mathbb{R} with the projection $p : x \mapsto x \bmod 1$.

3. Examples. We illustrate a few examples of simple coordination spaces and globally Pareto optimal path classes.

EXAMPLE 3.1. Consider the case where $N = 2$, $\Gamma_1 = \Gamma_2 = [-2, 2]$, and $\mathcal{O} = \{(x, y) : x^2 + y^2 < 1\}$, which corresponds to a pair of identical disc-shaped robots sliding along interval roadmaps which intersect in the workspace at right angles. There are exactly two globally Pareto optimal classes of paths from $(-2, -2)$ to $(2, 2)$, as illustrated in Fig. 3.1[left]. The difference between these two paths lies in which robot decides to pause in order to allow the other to pass through the intersection. Note also that this is a cylindrical coordination space and hence satisfies the Main Theorem above.




FIG. 3.1. The case of two disc-like robots translating along edges [left] yields a 2-d coordination space with a disc obstacle possessing two globally Pareto optimal classes [center]; Three robots with a spherical obstacle possesses a continuum of globally Pareto optimal classes [right].

EXAMPLE 3.2. We modify the previous example by letting $N = 3$ and choosing \mathcal{O} to be a round ball of radius 1 at the origin. By the symmetry of \mathcal{X} about the diagonal of the cube, it is clear that there is a circle’s worth of paths which begin at $(-2, -2)$, trace a straight line which is tangent to \mathcal{O} , and then exit this sphere tangentially with slope one. The projection of this family of paths to the first two coordinates includes as special cases the distinct Pareto optima of Example 3.1, as well as a continuum of paths whose goal times continuously interpolate the two. Hence, there is a continuum of globally Pareto optimal classes.

Note that the symmetry in these problems is not crucial: the number of globally Pareto optima in both examples is robust to small perturbations in \mathcal{O} . Note as well that the obstacle set in Example 3.2 is not generated by pairwise collisions of the robots; the obstacle set is determined by the positions of all three robots.

EXAMPLE 3.3. Coordination problems without the fixed path constraint can have continua of optima more easily than in the fixed path case. Consider the problem of translating two unit squares in the plane in such a way as to exchange their positions. Assume that the squares are centered at $(-1, 0)$ and $(1, 0)$. It is straightforward to exhibit a continuous family of distinct globally Pareto optimal classes. Extrema of this family consist of those coordinations for which one square translates exclusively in the horizontal direction, while the other square moves vertically in order to get out of the way. The intermediate coordinations involve one square translating up by an amount $0 < h < 1$ and the other translating down by a total amount $1 - h$. This example works with arbitrary translations, or with translations restricted to coordinate axis directions.

4. Discretization. Coordination spaces come equipped with a natural geometry as a subset of a (flat) product of metric graphs. For the remainder of this paper, we consider the PATH METRIC — the distance between points is defined in terms of the length of a shortest path (geodesic) between the points. Our tools for dealing with coordination spaces stem from knowledge of the geometry of the path metric.

4.1. Nonpositive curvature. Let (X, d) denote a complete metric space for which local geodesics exist: between any two sufficiently close points $p, q \in X$, there exists a path in X whose length is equal to $d(p, q)$. The notion of curvature bounds for such spaces is classical and can be found in the work of, e.g., Alexandrov (see [8] for references and a comprehensive introduction). We repeat that it is not necessary for X to be a manifold, much less smooth. In this ‘synthetic’ approach to geometry, one interrogates a space with geodesic triangles.

For each triple of points $p, q, r \in X$, draw the triangle in X with geodesics \overline{pq} , \overline{qr} , and \overline{rp} . Let p' , q' , and r' denote three points in the Euclidean plane \mathbb{E}^2 forming a COMPARISON TRIANGLE whose edges have length $d(p, q)$, $d(q, r)$, and $d(r, p)$ respectively. Choose a pair on point s and t on the edges of the geodesic triangle in X ; then, consider the corresponding points s' and t' in \mathbb{E}^2 . The CAT(0) INEQUALITY² for X is the following: for every p, q, r in X , and for every s, t , one has $d(s, t) \leq \|s' - t'\|$, where $\|\cdot\|$ denotes the Euclidean norm in \mathbb{E}^2 .

DEFINITION 4.1. *A complete geodesic space X is CAT(0) if the CAT(0) inequality holds for each geodesic triangle in X . A geodesic space X is NONPOSITIVELY CURVED or NPC if it is locally CAT(0) : that is, if the CAT(0) inequality holds for sufficiently small geodesic triangles in X .*

Otherwise said, X has nowhere positive curvature if small geodesic triangles in X are “no fatter” than similar triangles in the Euclidean plane. By assigning the appropriate definition of angles in a geodesic metric space [8], this condition translates into saying that the sum of the angles of a geodesic triangle in an NPC space is no larger than π . It is a standard fact that X is CAT(0) if and only if it is NPC and simply-connected [8].

²The term “CAT” is Gromov’s, used in deference to the work of Cartan, Alexandrov, and Toponogov. Alternately, one can take this as *comparare ab triangulos* (“to compare by triangles”).




FIG. 4.1. The coordination space of Example 3.1 is NPC [left] since triangles are “skinnier” than their Euclidean counterparts [center]; whereas that of Example 3.2 is not NPC [right].

The existence of *some* metric of nonpositive curvature on a space X implies several strong results about the topology and geometry of X . For example [8]: (1) the universal cover of X is contractible; (2) the fundamental group $\pi_1(X)$ has no torsional elements; and, most importantly for our purposes, (3) geodesics on X are unique up to homotopy, fixing the endpoints.

The problem of determining when a space is CAT(0) or NPC is surprisingly easy for piecewise-Euclidean cube complexes: that is, spaces build from Euclidean cubes glued together along faces. Since the individual cubes are flat, any nonzero curvature, if it exists in the complex, must be focused along the faces and detectable at the vertices of the cube complex. Intuitively, positive curvature arises whenever there is a convex “corner” in the complex.

More specifically, let v denote a vertex of a cube complex X . The number and types of cubes in X which have a corner at v are encoded in a combinatorial object called the LINK of v , see e.g. [8] for details. The following result of Gromov (which we have translated into simpler language for clarity) characterizes nonpositive curvature in terms of this structure, and is therefore called the LINK CONDITION:

THEOREM 4.2 (Gromov’s link condition). *A piecewise Euclidean cube complex is NPC if and only if the following holds for each vertex v . Given any set of $K \geq 3$ edges of X incident to v , any pair of which are part of a square in X , then all K edges are part of a K -dimensional cube with corner at v .*

This is illustrated for several 2-d complexes in Fig. 4.2. The intuition behind this theorem is that positive curvature in a cube complex is entirely determined by the presence of a ‘sharp’ corner.




FIG. 4.2. Vertices detect curvature for a cube complex: from left to right; positive, zero, and negative curvature in 2-d cube complexes. The link condition holds for the latter two of these complexes.

4.2. Discretized coordination spaces. The key step in our analysis of cylindrical coordination spaces is a spatial discretization of them into cubical complexes. Discretized configuration spaces have recently been used in other robotics contexts [1, 2, 16].

Consider a coordination space \mathcal{X} for roadmaps $\{\Gamma_i\}$. For simplicity, assume that all the edge lengths of the graphs Γ_i are rationally related. Because of this, one may choose a length $\delta > 0$ so that all edge lengths are integer multiples of δ . Denote by $\Gamma_i^{(n)}$ the graph obtained from Γ_i by subdividing it into edges of length $2^{-n}\delta$.

DEFINITION 4.3. *For each $n \geq 0$, denote by $\mathcal{X}^{(n)} \subset \Gamma_1^{(n)} \times \dots \times \Gamma_N^{(n)}$ the maximal closed subcomplex of the product cubical complex which does not intersect the obstacle set \mathcal{O} . This cubical complex is called the STAGE- n DISCRETIZED COORDINATION SPACE.*

The cubical nature of these spaces forms a convenient structure with which to manipulate paths. Our strategy is to approximate locally Pareto optimal paths with paths which follow along the diagonals of cubes in $\mathcal{X}^{(n)}$ for n sufficiently large. Working with these CUBE PATHS will allow us to skirt such technical issues as regularity of paths, etc. The following result is crucial:

THEOREM 4.4. *Cylindrical coordination spaces are NPC.*

Proof: By assumption, all coordination spaces in this paper are complete geodesic metric spaces under the path metric. We begin by showing that $\mathcal{X}^{(n)}$ converges to \mathcal{X} in the Hausdorff sense: i.e., given any point $p \in \mathcal{X}$ and an $\epsilon > 0$, there exists an $M > 0$ such that an ϵ -neighborhood of p intersects $\mathcal{X}^{(n)}$ for all $n > M$. This is certainly true for p in the interior of \mathcal{X} , since p is bounded away from \mathcal{O} . Since \mathcal{X} is collared, points in the boundary of \mathcal{X} are limits of points in the interior, and convergence follows.

It is known that a complete NPC space which is the Hausdorff limit of NPC spaces is itself NPC (see [8, Lemma II.3.10(1)], which is stated for CAT(0) spaces — apply this to the universal covers of \mathcal{X} and $\mathcal{X}^{(n)}$ to obtain the result). Hence, it suffices to show that $\mathcal{X}^{(n)}$ is NPC.

Since these are cubical complexes, Gromov’s link condition is necessary and sufficient to show NPC. Therefore, let v be any vertex of $\mathcal{X}^{(n)}$. Assume that there is a collection of $K \geq 3$ edges $\{e_i\}_1^K$ incident to v which pairwise bound squares in $\mathcal{X}^{(n)}$. Each e_i corresponds to an edge of $\Gamma_{\alpha(i)}$ for some indexing α satisfying $\alpha(i) \neq \alpha(j)$ for all $i \neq j$. The square bound by edges e_j and e_k corresponds to a square in

$$\Gamma_{\alpha(j)} \times \Gamma_{\alpha(k)} - \Delta_{\alpha(j),\alpha(k)}$$

for each j and k . By the cylindrical structure of the holes in \mathcal{X} , the product of the edges $\{e_i\}$ is a K -dimensional cube in $\prod_1^m \Gamma_{\alpha(i)}$ which does not intersect $\Delta_{\alpha(j),\alpha(k)}$ for any pair j, k . This therefore defines a K -dimensional cube in $\mathcal{X}^{(n)}$ spanning the edges. The link condition is satisfied. \diamond

Theorem 4.4 generalizes a number of known results about NPC configuration spaces, especially the results of [1] about configuration spaces of points on graphs, and more general “colored” configuration spaces of graphs.³

³This proof is simpler than that of Abrams in that it requires no direction computation of the homotopy groups of the Hausdorff limit.

4.3. Cube paths. The spatial discretization of coordination spaces extends to a discrete version of paths. The following notion of discrete paths is crucial to future proofs.

DEFINITION 4.5. Let X denote a cubical complex. A CUBE PATH from vertices v_0 to v_N in X is an ordered sequence of closed cubes $\mathcal{C} = \{C_i\}_1^N$ of X which satisfy (1) $C_i \cap C_{i+1} = \{v_i\}$, a vertex; and (2) C_i is the smallest cube of X containing v_{i-1} and v_i . A cube path is said to be NORMAL if in addition $\forall i$ (3) $C_{i+1} \cap \text{St}(C_i) = v_i$, where $\text{St}(C_i)$ is the STAR of C_i (all cubes in X , including C_i , which have C_i as a face).

To each cube path is associated a PL (piecewise linear) path in X given by the chain of straight line segments from v_i to v_{i+1} for $i = 0 \dots N - 1$. Roughly speaking, a normal cube path is one which uses the highest dimensional cubes as early as possible in the path.⁴ This hints at the intuition of normal paths as locally Pareto optimal, an intuition which we demonstrate is entirely justified. The key fact about normal cube paths is that, like geodesics on an NPC space, they are unique up to homotopy:




FIG. 4.3. Examples of normal cube paths (with the diagonals highlighted).

COROLLARY 4.6. Any cube path between vertices of $\mathcal{X}^{(n)}$ is homotopic to a unique normal cube path.

Proof: This is true for any NPC cube complex via Prop. 3.3 of [26]. By Theorem 4.4, this holds for $\mathcal{X}^{(n)}$. \diamond

5. Topological bounds on Pareto optima. In this section we demonstrate a uniqueness result for locally Pareto optimal classes. The strategy of our proof is to show that any locally Pareto optimal path class in \mathcal{X} contains a representative which is a limit of normal cube paths in $\mathcal{X}^{(n)}$. For the remainder of the paper, \mathcal{X} will denote a cylindrical roadmap coordination space.

The geometry on \mathcal{X} is that defined by the path metric: the metric distance between points of \mathcal{X} is defined to be the length of a geodesic path between the points. Unless otherwise noted, the default notion of length is the standard length obtained by integrating the ℓ^2 norm of the velocity vectors over the path. Other notions of length are useful. Given a velocity vector, the ℓ^∞ norm of that vector is the usual definition: in this case, the maximal speed of the AGV's is traveling at that instant. The ℓ^∞ length of a path in \mathcal{X} is the integral of the ℓ^∞ norm of the velocity vectors over the path. This ℓ^∞ length measures the total elapsed time that the AGV's take to execute the coordination implied by the path. As such,

⁴These paths are called “normal” in [26] since they are used to generate normal forms for languages of fundamental groups acting on CAT(0) cube complexes.

ℓ^∞ geodesics between points correspond to “fastest” coordinations.

LEMMA 5.1. *Any ℓ^∞ geodesic path on \mathcal{X} can be C^0 - and ℓ^∞ -approximated by a cube path in $\mathcal{X}^{(n)}$ for n sufficiently large.*

Proof: Given an ℓ^∞ geodesic α , those portions of the curve which lie near $\partial\mathcal{O}$ may not be in $\mathcal{X}^{(n)}$. However, given the embedded collar $\Phi : \partial\mathcal{O} \times [0, 1] \hookrightarrow \mathcal{X}$ from Assumption 2.2, the curve α can be pushed in to the interior of \mathcal{X} by increasing the second coordinate of Φ . This yields a curve β which avoids $\partial\mathcal{O}$, is freely homotopic to α , and is both C^0 and ℓ^∞ -close.

We next approximate β by a cube path. Since the image of β is compact, there exists an $\epsilon > 0$ so that a tube of radius ϵ about the image of β does not intersect $\partial\mathcal{O}$. Choose n so that the grid size $\delta = 2^{-n}$ is less than ϵ . Modify β so that its endpoints are at vertices of $\mathcal{X}^{(n)}$. Consider the vertex $v^0 = \beta(0)$ and the ℓ^∞ balls $B_r = B_r^\infty(v^0)$ of radius r about v^0 in $\mathcal{X}^{(n)}$. Let $w^1 = \beta \cap \partial B_\delta$ and $w^2 = \beta \cap \partial B_{2\delta}$.




FIG. 5.1. *Approximation of a path by a cube path.*

Let $F = B_\delta(v^0) \cap B_\delta(w^2)$. Clearly, $w^1 \in F$ and for any $x \in F$ it holds that the PL path $v^0 \rightarrow x \rightarrow w^2$ is of the same ℓ^∞ length as the restriction of β to this neighborhood. We claim that F contains a vertex of the grid structure on $\mathcal{X}^{(n)}$. To see this, note that F is a subset of some face of B_δ defined by the set of all points $x = (x_i)$ such that

$$x_i \in [v_i^0 - \delta, v_i^0 + \delta] \cap [w_i^2 - \delta, w_i^2 + \delta] \quad ; \quad i = 1 \dots N \quad (5.1)$$

Since these intervals have the same length, it follows that either $v_i^0 - \delta$ or $v_i^0 + \delta$ lies within the intersection for each i . Therefore, F contains a gridpoint v^1 of $\mathcal{X}^{(n)}$, and we may replace the segment of the path β with a PL path from v^0 to v^1 to w^2 . Repeat the argument inductively, beginning at the vertex v^1 and considering the ℓ^∞ balls about v^1 of radius δ and 2δ .

Marching along yields a PL path passing through vertices of $\mathcal{X}^{(n)}$ which are sequentially of ℓ^∞ -distance δ . This is therefore a cube path in $\mathcal{X}^{(n)}$. This completes the proof of the lemma, as all modifications are homotopies within a neighborhood of the initial path α . \diamond

LEMMA 5.2. *Let α be a locally Pareto optimal path on \mathcal{X} . One can find homotopic cube paths with endpoints close to those of α such that all the goal times approximate those of α .*

Proof: Any locally Pareto optimal path is a concatenation of paths which go from the attainment of the i^{th} goal to the $(i + 1)^{\text{st}}$ goal. Each of these subpaths is an ℓ^∞ geodesic

Algorithm 1 LeftGreedy(\mathcal{C})

Require: $\mathcal{C} = \{C_i\}_1^m$ is an ℓ^∞ -geodesic cube path

- 1: **while** \mathcal{C} is not normal **do**
 - 2: **for** $i = 1..m$ **do**
 - 3: $Z \leftarrow C_{i+1} \cap \text{St}(C_i)$
 - 4: $C_{i+1} \leftarrow C_{i+1}/Z$
 - 5: $C_i \leftarrow C_i \times Z$
 - 6: **end for**
 - 7: **end while**
-

whose ℓ^∞ -length contributes to the elapsed time. Applying Lemma 5.1 inductively to these subpaths gives a path which approximates all the elapsed times. \diamond

A simple algorithm, **LeftGreedy**, converts a cube path \mathcal{C} into a normal cube path by sweeping along \mathcal{C} from left to right, comparing the star of each cube with its neighbor to the right (line 3), and shifting the common directions to the left cube (lines 4-5).

LEMMA 5.3. *Given a locally Pareto optimal cube path \mathcal{C} , **LeftGreedy**(\mathcal{C}) is a normal cube path which is Pareto equivalent to \mathcal{C} .*

Proof: Since common directions are shifted to the left, each sweep through \mathcal{C} for which $C_{i+1} \cap \text{St}(C_i) \neq \{v_i\}$ for some i strictly decreases the positive integer-valued function $\sum_i i \cdot \dim(C_i)$; hence, convergence.

To argue that **LeftGreedy** can be executed through Pareto equivalence, note first that the shifting process of lines 4-5 results in a homotopic path, since all modifications take place within $\text{St}(C_i) \cup C_{i+1}$. If the j^{th} coordinate x_j reaches its goal at vertex v_i of \mathcal{C} , then all subsequent vertices v_k ($k \geq i$) lie within the hyperplane of \mathcal{X} having the same j^{th} coordinate. Hence, in line 3, the common factor Z is trivial in the j^{th} coordinate for v_i and beyond. Thus, the time-to-goal for x_j is unchanged by steps 4-5. \diamond

We now may prove the Main Theorem:

THEOREM 5.4. *Locally Pareto optimal path classes on \mathcal{X} are in bijective correspondence with homotopy classes of paths (fixing the endpoints).*

Proof: Existence of a locally Pareto optimum does not rely on the cylindrical nature of \mathcal{X} , but rather on the fact that it is compact and piecewise real-analytic. Via Lemma 2.6 applied to each homotopy class, the existence of a locally Pareto optimal path is guaranteed by the existence of a path which minimizes some monotone scalarization of the cost function, e.g., average time-to-goal. For fixed endpoints, this scalar function is positive and bounded below over the space of piecewise real-analytic paths in a given homotopy class.

The classification theorem relies on the cylindrical nature of \mathcal{X} . Assume there are two locally Pareto optimal paths, α and α' , which are homotopic. Approximate these paths by cube paths \mathcal{C} and \mathcal{C}' on $\mathcal{X}^{(n)}$ for some large n via Lemma 5.2. The approximations \mathcal{C} and \mathcal{C}' are homotopic in \mathcal{X} to α and α' respectively: we claim that \mathcal{C} and \mathcal{C}' are homotopic in $\mathcal{X}^{(n)}$. Since α and α' are homotopic in \mathcal{X} , there exists a map of a closed disc into \mathcal{X} whose boundary is sent to $\alpha * (\alpha')^{-1}$ (that is, the loop obtained by concatenating loop α with the

reverse of the loop α'). The collar Φ from Assumption 2.2 which is used to push α and α' into the interior of $\mathcal{X}^{(n)}$ also pushes this spanning disc into the interior. Thus, the paths \mathcal{C} and \mathcal{C}' are homotopic within $\mathcal{X}^{(n)}$.

Feed these cube paths into the algorithm `LeftGreedy`. Via Lemma 5.3, the outputs are normal cube paths which are Pareto equivalent to \mathcal{C} and \mathcal{C}' respectively. By Corollary 4.6, \mathcal{C} and \mathcal{C}' are the same path and hence have the same output times. Therefore, α and α' have output times which are arbitrarily close. \diamond

This implies that the locally Pareto optimal classes are “discrete” and cannot arise in continua as in the non-cylindrical case. It does not automatically follow that there is a bound on the number of globally Pareto optimal paths, since any cylindrical coordination space \mathcal{X} which is not simply connected must have a countably infinite number of homotopy classes of paths. However, since the obstacle set is collared, the fundamental group of \mathcal{X} is finitely generated, and, apart from this finite collection of generators, all other paths “wrap around” a loop multiple times. One rightly suspects that such repetition makes everyone’s arrival times later, and thus cannot be a globally Pareto optimal solution. We make this intuition rigorous below.

THEOREM 5.5. *The number of globally Pareto optimal paths between fixed endpoints on \mathcal{X} is finite.*

Proof: Fix endpoints and consider $\{\gamma_i\}_1^\infty$ an infinite collection of non-equivalent globally Pareto optimal paths. By Theorem 5.4, all these are in different homotopy classes. The endpoints of these paths in the universal cover must be an unbounded discrete set, whose ℓ^∞ lengths tend uniformly to infinity, after taking a subsequence and reordering.

Denote by $\{T_i^j\}_{i=1..N}^{j=1..N}$ the $j = 1..N$ goal times of the path γ_i . By the pidgeonhole principle, one may assume (after taking a subsequence) that for some k , $T_i^{k+1} \geq T_i^k + i$ for all i . By compactness of \mathcal{X} , the set of points $\{\gamma_i(T_i^k)\}$ has an accumulation point x_k . Likewise, the set of points $\{\gamma_i(T_i^{k+1})\}$ has an accumulation point x_{k+1} . For all i large, a small perturbation to the path γ_i after T_i^k in a neighborhood of x_k and x_{k+1} reduces T_i^j by more than $i - 1$ for each $j > k$, without increasing any other goal times. This contradicts the assumption of globally Pareto optimal paths. \diamond

COROLLARY 5.6. *There exists a finite bound on the number of globally Pareto optimal classes on \mathcal{X} between fixed endpoints which is independent of the endpoints.*

Proof: For each fixed pair of endpoints in $\mathcal{X} \times \mathcal{X}$, there is a finite number of globally Pareto optimal classes via Theorem 5.5. The bound extends to small open neighborhoods in $\mathcal{X} \times \mathcal{X}$ via the same argument. Thanks to the compactness of \mathcal{X} , this open cover restricts to a finite subcover. \diamond

This completes the proof of the Main Theorem.

6. Computing locally Pareto optima. The proof of the Main Theorem reveals a technique for computing “canonical” representatives for the locally Pareto optimal classes. One observes from the proof that passing to a sufficiently fine cubical discretization of the coordination space and choosing a normal cube path between endpoints yields an approxi-

mation to a locally Pareto optimal path. Normal cube paths are not difficult to compute, given the cubical approximation — one simply cuts across as many diagonals as early as possible in the path.

However, this procedure is not optimal as a means of generating representatives of a Pareto equivalence class, since it involves a potentially large and expensive discretization step and it yields only an approximation to the desired optimum. A less expensive approach that yields the true locally Pareto optimal representative path comes from deriving a ‘continuous’ version of a normal cube path which does not rely on having a cubical structure at all. These paths, called LEFT-GREEDY NORMAL paths, are well-defined, unique up to homotopy, and locally Pareto optimal.

A companion paper to this [17] gives proper definitions of left-greedy normal paths, along with algorithms for their computation, complexity bounds, and implementation in the context of coordinating several irregular-shaped robots translated along planar tracks. It is an interesting feature that the algorithms for computing these locally Pareto optimal paths rely crucially on the nonpositive curvature of the underlying coordination space.

7. Moving obstacles. Assume that, as before, we have N robots which traverse roadmaps $\{\Gamma_i\}_1^N$. If, as will typically occur, the obstacle set \mathcal{O} is defined by pairwise collisions, then the resulting coordination space \mathcal{X} will be cylindrical and the finiteness bounds apply.

Assume now that there are additional objects which, like the robots, move through the shared environment, but which, unlike the robots, are not controllable. These objects move along pre-determined trajectories in the workspace. We make no assumptions about their speed, shape, or motion.⁵ We assume *only* that the trajectory through the workspace is a fixed function of time. Such a system gives rise to a roadmap coordination system with moving obstacles.

We consider the problem of determining locally Pareto optimal coordinations of the N robots through the workspace, where all admissible paths are those which avoid collision with the moving obstacles. Our strategy is to augment the coordination space by adding one dimension for time, similar to the methods used in [20, 30].

THEOREM 7.1. *Any roadmap coordination system with moving obstacles possesses a finite number of globally Pareto optimal classes of paths.*

Proof: Let \mathcal{X} denote the coordination space of the N robots obtained by ignoring the additional obstacles. We will build an augmented coordination space of the form $\overline{\mathcal{X}} \subset [0, T] \times \left(\prod_{k=1}^N \Gamma_k\right)$. Encode the positions of the moving objects as a time-parameterized subset of the workspace (where, without a loss of generality, $0 \leq t \leq T$ for some T). For each $1 \leq i \leq N$, denote by $\overline{\mathcal{O}}_i \subset [0, T] \times \Gamma_i$ the set of points of the form (t, x_i) where the i^{th} robot in configuration x_i is in collision with the moving obstacles’ positions at time t . We define $\overline{\mathcal{X}}$ via the obstacle set:

$$\overline{\mathcal{O}} = ([0, T] \times \mathcal{O}) \bigcup_i \left\{ (t, x_1, \dots, x_N) \in [0, T] \times \prod_{k=1}^N \Gamma_k : (t, x_i) \in \overline{\mathcal{O}}_i \right\}. \quad (7.1)$$

Claim 1: The augmented coordination space $\overline{\mathcal{X}}$ is cylindrical.

⁵Indeed, the moving obstacles may, during their trajectory, change shape, speed, or even connectivity.

To prove this, note that $\overline{\mathcal{X}}$ is obtained from $[0, T] \times \prod_i \Gamma_i$ by removing sets defined either from \mathcal{O} or from \mathcal{O}_i . These are all pairwise defined cylindrical subsets.

End of Claim 1.

We now apply the same proof to $\overline{\mathcal{X}}$ with two key modifications. The first is that the time direction $[0, T]$ in $\overline{\mathcal{X}}$ is rigid: all admissible paths in $\overline{\mathcal{X}}$ must travel with unit speed in the t -direction. The second is that the endpoint is variable: paths start at the point $(0, p)$ and end at (T', q) for some fixed $p, q \in \mathcal{X}$ and for some variable $T' \leq T$.

For the moment, assume that T' is fixed.

Claim 2: The number of admissible globally Pareto optimal paths from $(0, p)$ to (T', q) is finite.

To prove this, one shows that the proofs of Theorems 5.4 and 5.5 apply to admissible paths (those with unit speed in the t -direction). Not all homotopy classes of $\overline{\mathcal{X}}$ are represented, but only those which have unit speed in the t -direction. The crucial observation for the analogue of Theorem 5.4 is that Algorithm 1 applied to a path with unit speed in the t -direction preserves this property. The analogue of Theorem 5.5 in this setting is unchanged.

End of Claim 2.

In the general case, the endpoint has t -coordinate T' . For each such T' there is a finite set of globally Pareto optimal paths in $\overline{\mathcal{X}}$.

Claim 3: The projections of these globally Pareto optimal paths in $\overline{\mathcal{X}}$ to \mathcal{X} is independent of T' , for T' sufficiently large.

To prove this, assume that \mathcal{C}' is a normal cube path from $(0, p)$ to (T', q) in $\overline{\mathcal{X}}$. Let \mathcal{C} denote the cube path obtained by appending to \mathcal{C}' the straight line segment from (T', q) to (T, q) . Claim 3 follows if \mathcal{C} is also a normal cube path.

Assume that \mathcal{C} is not normal. Then, since both pieces of \mathcal{C} are normal (\mathcal{C}' and the segment $(T', q) \rightarrow (T, q)$), the normal condition must fail precisely at the vertex (T', q) . This would imply that the star of the last cube in \mathcal{C}' intersects the segment $(T', q) \rightarrow (T, q)$; hence, the last cube of \mathcal{C}' had no component in the t -direction. However, this is a contradiction, since \mathcal{C}' is an admissible path with constant speed 1 in the t -direction.

End of Claim 3 and the Theorem. ◇

The problem of computing locally Pareto optimal paths is no different in this case than in the stationary obstacle case of the previous sections. The motion of the obstacles is incorporated in the geometry of the coordination space in the time-dimension.

8. On positive curvature. We claim that the crucial reason for the finiteness bound is that the cylindrical coordination spaces can be approximated by cubical complexes without any vertices of positive curvature. For more general coordination spaces, all the steps of the proof of the finiteness bound hold *except* the uniqueness of normal cube paths up to homotopy.

For example, if one approximates the space of Example 3.2 by discretizations $\mathcal{X}^{(n)}$, then, as before, there are a finite number of normal cube paths. For example, in $\mathcal{X}^{(2)}$, \mathcal{O} is approximated by a cube and there are exactly three normal cube paths. However, they

are not unique up to homotopy, and the number of such paths grows exponentially in the discretization step n , giving rise to the continuum of optima in the limit: see Fig. 8.1




FIG. 8.1. As one approximates the obstacle of Example 3.2 by finer and finer cubes, the number of normal cube paths from one corner to its antipode grows exponentially.

The positive curvature in Example 3.2 is transparent: the obstacle set \mathcal{O} clearly has a positively curved boundary. But in Example 3.3, the positive curvature is more subtle and arises from the fact that the obstacle set \mathcal{O} is not axis-aligned. Consider first the case in which the squares have full mobility in all directions. We discretize this \mathcal{X} to a cube complex by letting edges correspond to sliding a square in either the horizontal or vertical direction. Higher dimensional cubes correspond to slides which “commute” (they are compatible without causing a collision between the two squares).

We claim that positive curvature is created in the discretization. It suffices to demonstrate the failure of the link condition: namely, to show that there exists a triple of edges which pairwise commute but which together yield an illegal state. This is illustrated in Fig. 8.2. Each pair of the three moves exhibited there commutes (with one pair generating a diagonal slide), but all three performed simultaneously generates a collision.




FIG. 8.2. The link condition fails for cubical approximations to the coordination space of Example 3.3: the three translations shown can be performed simultaneously pairwise, but performing all three at once causes a collision.

There is one last type of failure on which we remark. If we modify Example 3.3 so that squares are constrained to slide solely in the horizontal and vertical directions, then, indeed, the discretized cube complex now has no positive curvature, since sliding the same square horizontally and vertically are no longer commutative moves. Because of this, however, there are nontrivial loops in $\pi_1(\mathcal{X})$ generated by moving a single square north-east-south-west in

turn. There is one such loop for each vertex in $\mathcal{X}^{(n)}$, and the length of the loop is four times the discretized edge length. As $n \rightarrow \infty$, the number of such loops becomes arbitrarily large while their lengths become arbitrarily small: there is a “froth” of π_1 . Even though normal cube paths (and hence locally Pareto optimal representatives) are unique up to homotopy, the number of such homotopy classes blows up exponentially. In this case, the argument involving π_1 being finitely generated no longer holds and, indeed, there are an uncountably infinite number of globally Pareto optimal paths in the limit.

The techniques we have introduced here — CAT(0) geometry and nonpositive curvature — are very classical concepts. Yet their applications outside of geometry and algebra are in the very beginning stages, and we are hopeful that other problems in robotics and optimization will benefit from these perspectives. We end by noting that the problem of computing global Pareto optima is not dissimilar to that of computing Euclidean geodesics. It is well-known that computing shortest paths in 3-d is NP-hard in general [10] (see, *e.g.*, [24] for an up-to-date account of hardness results for geodesics). It is worth noting that every example surveyed in [24] which has an NP-hard geodesic problem is not a CAT(0) space. It would be a good omen for the applicability of CAT(0) techniques if the geodesic problem simplified in the case of CAT(0) spaces in general. In a parallel application, pursuit-evasion problems on domains of dimension greater than two depend crucially on the geometry and topology of the domain. For certain pursuit games, the existence of a successful pursuit strategy can be guaranteed on a CAT(0) space [4].

REFERENCES

- [1] A. Abrams. *Configuration spaces and braid groups of graphs*. Ph.D. thesis, UC Berkeley, 2000.
- [2] A. Abrams and R. Ghrist. State complexes for metamorphic robot systems. *Intl. J. Robotics Research*, **23**(7,8), 809–824.
- [3] 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.
- [4] S. Alexander, R. Bishop, and R. Ghrist. Pursuit and evasion in arbitrary dimensions. Submitted.
- [5] M. D. Ardema and J. M. Skowronski. Dynamic game applied to coordination control of two arm robotic system. In R. P. Hämmäläinen and H. K. Ehtamo, editors, *Differential Games - Developments in Modelling and Computation*, pages 118–130. Springer-Verlag, Berlin, 1991.
- [6] J. Barraquand and J.-C. Latombe. Robot motion planning: A distributed representation approach. *Int. J. Robot. Res.*, 10(6):628–649, December 1991.
- [7] Z. Bien and J. Lee. A minimum-time trajectory planning method for two robots. *IEEE Trans. Robot. & Autom.*, 8(3):414–418, June 1992.
- [8] M. Bridson and A. Haefliger, *Metric Spaces of Nonpositive Curvature*, Springer-Verlag, Berlin, 1999.
- [9] S. J. Buckley. Fast motion planning for multiple moving robots. In *IEEE Int. Conf. Robot. & Autom.*, pages 322–326, 1989.
- [10] J. Canny and J. Reif. Lower bounds for shortest path and related problems. In *Proc. 28th Ann. IEEE Symp. Found. Comp. Sci.*, 49–60, 1987.
- [11] 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.
- [12] B. Chazelle, H. Edelsbrunner, M. Grigni, L. J. Guibas, J. Hershberger, M. Sharir, and J. Snoeyink. Ray shooting in polygons using geodesic triangulations. *Algorithmica*, 12:54–68, 1994.
- [13] 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.
- [14] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer, Berlin, 1997.
- [15] M. Erdmann and T. Lozano-Perez. On multiple moving objects. In *IEEE Int. Conf. Robot. & Autom.*, pages 1419–1424, 1986.
- [16] R. Ghrist. Shape complexes for metamorphic robot systems. In *Algorithmic Foundations of Robotics*

- V, *STAR 7*, pages 185–201, 2004.
- [17] R. Ghrist, J. O’Kane, and S. M. LaValle. Pareto optimal coordination on roadmaps. *Intl. J. Robotics Research*, 12(11), 997–.
 - [18] A. Hatcher. *Algebraic Topology*. Cambridge University Press, 2001.
 - [19] 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.
 - [20] K. Kant and S. W. Zucker. Toward efficient trajectory planning: the path-velocity decomposition. *Intl. J. Robotics Research* 5(3), 72–89, 1986.
 - [21] D.G. Kirkpatrick. Optimal search in planar subdivisions. *SIAM J. Comput.*, 12(1):28–35, 1983.
 - [22] 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. & Autom.*, pages 1847–1852, May 1994.
 - [23] 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.
 - [24] J. S. B. Mitchell and M. Sharir. New results on shortest paths in three dimensions. In *Proc. 20th Annual ACM Symposium on Computational Geometry*, 124–133, June 2004.
 - [25] 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.
 - [26] G. A. Niblo and L. D. Reeves. The geometry of cube complexes and the complexity of their fundamental groups. *Topology*, 37(3):621–633, 1998.
 - [27] 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.
 - [28] V. Pareto. *Cours d’Économie Politique*, Lausanne, 1896.
 - [29] L. E. Parker. Cooperative motion control for multi-target observation. In *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, pages 1591–1598, 1998.
 - [30] 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.
 - [31] Y. Sawaragi, H. Nakayama, and T. Tanino. *Theory of Multiobjective Optimization*. Academic Press, New York, NY, 1985.
 - [32] 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.
 - [33] 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.
 - [34] 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.
 - [35] 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.