Maximum Thick Paths in Static and Dynamic Environments

Esther M. Arkin
estie@ams.sunysb.edu
Deptartment of Applied Mathematics and Statistics
Stony Brook University
Stony Brook, NY 11794-3600
USA

Joseph S. B. Mitchell
jsbm@ams.sunysb.edu
Department of Applied Mathematics and Statistics
Stony Brook University
Stony Brook, NY 11794-3600
USA

Valentin Polishchuk
valentin.polishchuk@helsinki.fi
Helsinki Institute for Information Technology
University of Helsinki and Helsinki University of Technology
Finland

ABSTRACT

We consider the problem of finding a maximum number of disjoint paths for unit disks moving amidst static or dynamic obstacles. For the static case we give efficient exact algorithms, based on adapting the "continuous uppermost path" paradigm. As a by-product, we establish a continuous analogue of Menger's Theorem. (In this extended abstract we only state these results.)

Next we study the dynamic problem in which the obstacles may move, appear and disappear, and otherwise change with time in a known manner; in addition, the disks are required to enter/exit the domain during prescribed time intervals. We observe that (unless P=NP), for any $\alpha, \beta > 0$, one cannot decide in polynomial time whether there exist $\lceil \alpha K \rceil$ paths for disks of radius βR, where K is the maximum number of paths for radius-R disks. The problem is hard even if the obstacles are static, and only the entry/exit time intervals are specified for the disks. This motivates studying "dual" approximations, compromising on the radius of the disks and on the maximum speed of motion.

Our main result is a pseudopolynomial-time dual-approximation algorithm: if K unit disks, each with unit bound on the speed, may be routed through an environment, our algorithm finds (at least) K paths for disks of radius $\Omega(1)$ moving with speed $O(1)$. The algorithm computes a maxflow with "forbidden pairs" in an "adaptive" grid, laid out in space-time. Although (as we show) in general finding even an approximation to the maxflow with forbidden pairs is not possible (unless P=NP), a careful choice of time discetization and a non-uniform grid of "way-points" allows us to give provable approximation guarantees on the quality of the solution produced by the algorithm.

Our algorithm extends to higher dimensions and to finding paths for translational motion of arbitrary-shape objects.

Categories and Subject Descriptors

F.2.2 [**Analysis of Algorithms and Problem Complexity**]: Nonnumerical Algorithms and Problems—*Geometrical problems and computations*

General Terms

Algorithms

Keywords

Motion planning, approximation algorithms

SCG'08, June 9–11, 2008, College Park, Maryland, USA.

1. INTRODUCTION

Path planning in geometric domains is an important computational geometry subject with applications in robotics, VLSI routing, air traffic management (ATM), sensor networks, etc. In many applications it is of interest to find *multiple* disjoint paths for *non-point* objects avoiding *moving* obstacles. This is the problem studied in this paper.

The input to the problem is specified by a polygonal domain, with two edges of the outer polygon designated as the "source" and the "sink". The holes/obstacles in the domain move along known trajectories. Also given are the entry and exit time intervals. The goal is to find a maximum number of trajectories for equal-radii disks, moving with bounded speed, entering (resp., exiting) the domain through the source (resp., sink) during the entry (resp., exit) interval, never intersecting each other, nor the obstacles.

Specifying entry/exit times is not crucial for the problem's (in)tractability: one may just introduce obstacles that cover the source/sink during the times that are infeasible for entry/exit. We opted to retain the entry/exit time constraints in our formulation for two reasons. First, our algorithm is capable of addressing them; moreover, it can also respect given bounds on the total number of disks entering/exiting the domain within any set of *multiple* time intervals. This is important in ATM—our motivating application—since it allows us to bound the workload of air traffic controllers for a given sector of airspace during any specified set of time periods. Second, ATM often involves putting aircraft into "holding patterns" due to congestion at an airport or in an adjacent sector. Even if an aircraft may pass through an airspace quickly, it may be necessary to introduce delay (through path maneuvers or speed control). Monitoring aircraft in holding patterns adds significantly to the workload for a human controller; it is preferable to monitor aircraft that move as directly as possible through an airspace sector.

Related Work

Finding a maximum number of source-sink paths in *graphs* is equivalent to computing maxflow. Three classical graph-theoretic results related to the equivalence are the MaxFlow-MinCut, the Flow Decomposition, and Menger's Theorem. The extensions of the theorems to (static) geometric domains are developed in [13, 15, 22], and this paper, respectively. Real-data implementation of the algorithm, based on our Continuous Menger's Theorem, is reported in [8, 9, 16].

A lot of research has been done on computing *fastest* paths amidst moving/morphing obstacles for a point object, possibly with nonholonomic motion constraints, both in 2D and 3D [1, 2, 5–7, 17, 21, 23–25, 27–29]. See the books [11, 12], Section 4.4 in the survey [14], and references therein for details.

To the best of our knowledge, no algorithm with *provable* approximation guarantees has been known previously for finding a maximum number of paths in the presence of moving obstacles. The approaches for planning paths for multiple objects fall into two major categories: *prioritized*, when the paths are routed one-by-one, and *coordinated*, ranging from centralized to roadmap-based to decoupled (the classification is taken from [26], where a more detailed discussion of the approaches may be found). Existing algorithms are *heuristics*, many of them are based on laying a regular grid in time-space and searching it for paths. In ATM applications, the possibly leading methods are the heuristics implemented in the Flow-Based Route Planner (FBRP) [19]. The FBRP searches the grid for the paths greedily and iteratively: each computed (thick) path becomes a constraint (obstacle) in space-time for subsequent paths. Although examples exist for which arbitrarily many paths may be routed, while the FBRP produces only one, the planner performs very well in practical situations [8–10, 19, 20].

Our Contribution

Similar to existing heuristics, we use a uniform discretization of time. The novelty of our solution is in the "discretization" of the space: instead of using a regular grid, we pack, at each time slice, a maximal number of disks in the domain. We then find maxflow in a graph built on the disks; the flow decomposes into a set of paths. We prove that the conflicts between the paths, introduced by the discretization, can be resolved locally. We show how to balance the time discretization step, the radius of the disks packed, and the disks' speed: we prove that if there exist K paths through the domain for unit disks moving with speed at most 1, our algorithm will find, for any $\Delta t < 1/2$, (at least) K paths for radius-$(\frac{1}{3} - \frac{1}{2}\Delta t)^2$ disks moving with speed at most $10/\Delta t$. We give hardness results that justify the need for approximate solutions.

For the case of static obstacles, we give exact polynomial-time algorithms based on a modification of continuous-Dijkstra-type uppermost shortest path algorithm for maxflow in a polygonal domain [13]. As a by-product, we formulate and prove the Continuous Menger's Theorem — an extension of the famous graph theorem to geometric domains.

Note that the number of paths that exist in a domain may be exponential in the input size (e.g., there may exist $\Omega(N)$ paths in a $2 \times N$ rectangle — specified with $O(\log N)$ bits). Nevertheless, for the case of static obstacles we can, in strongly polynomial time, output a succinct representation of the paths, based on the Continuous Flow Decomposition Theorem [15]. Our algorithm for the case of moving obstacles runs in pseudopolynomial time.

2. STATIC OBSTACLES

The input to our problem is a polygonal domain Ω defined by an outer (simple) polygon P and a set $\mathcal{H}$ of holes in it. Two edges, Γ_I and Γ_O, of P are designated as the *source* and the *sink*. A *thick path* is the Minkowski sum of a curve and unit disk. We find the maximum number of thick disjoint Γ_I-Γ_O paths in Ω by modifying the continuous-Dijk-stra-type uppermost shortest path algorithm for finding maxflow in polyhedral domains [13], taking into account the discrete nature of our problem. In this abstract, we summarize the statements of our results; see [18, Section 6.3] for the details.

THEOREM 2.1. *Let n be the number of vertices of Ω, let $h = |\mathcal{H}|$ be the number of holes in it. A representation of the maximum number of thick non-crossing Γ_I-Γ_O paths can be found in $O(nh + n \log n)$ time. If the paths have different thicknesses, the problem is NP-hard unless the* order *of the paths along Γ_I/Γ_O is specified, in which case our algorithms apply. In a simple polygon the problems can be solved in linear time.*

THEOREM 2.2. **Continuous Menger's Theorem.** *Let B (resp., T) be the part of ∂P counterclockwise between Γ_I and Γ_O (resp., Γ_O and Γ_I). The maximum number of thick non-crossing paths in Ω is equal to the length of a shortest T-B path in the "thresholded critical graph", in which there is a node for each hole, for T, and for B, and the length of an edge is the floor of the distance between the holes (or B or T), divided by 2.*

3. MOVING OBSTACLES

In this section we consider the dynamic version of the problem: Given a polygonal domain with *moving* obstacles, find a maximum number of paths for unit disks moving with bounded speed, each path going from the source to the sink, so that no disk ever intersects an obstacle and two disks never collide. We observe that (unless P=NP) no polynomial-time algorithm can be given for the problem, even if arbitrary compromise on the number of the paths and the radius of the disks is allowed. We also prove that the problem is NP-hard even if the obstacles are static, but the disks must enter and exit the domain during prescribed time intervals. We then show how to find the optimal number of paths for smaller disks that are allowed to break the speed limit slightly.

The idea of our solution is as follows. We prove that, when lifted to the (x, y, t)-space, every feasible path contains a "stack" of cylinders; the consecutive cylinders in the stack overlap a lot by height, and are only slightly shifted horizontally — this allows one to find a chain of *oblique* cylinders inside the stack. We then search the graph, built from the oblique cylinders, for a maximum number of disjoint paths; although the paths found may self-intersect, we show how to "bend" them locally to resolve the intersections. The algorithm is presented in Fig. 5; in the reminder of the section we describe its details and prove its correctness.

As before, let $\Omega = (P, \mathcal{H})$, Γ_I, Γ_O, denote the domain, the source and the sink. The positions and shapes of the holes/obstacles are now functions of time, $\mathcal{H} = \mathcal{H}(t)$; let

$\Omega(t) = (P, \mathcal{H}(t))$. Let T_I and T_O be the *entry and exit time intervals*. Assume, w.l.o.g., that $\min\{t|t \in T_I\} = 0$, and let $T = \max\{t|t \in T_O\}$; the interval $[0, T]$ is called the *planning horizon*. A *path* π is now a continuous curve in the domain, $\pi : [t_I^\pi, t_O^\pi] \mapsto \Omega$, $t_I^\pi \in T_I$, $t_O^\pi \in T_O$, parameterized by time, $\pi = \pi(t)$. For $r > 0$ and $S \subset \mathbb{R}^2$ let $< S >^{(r)}$ denote the Minkowski sum of S and the disk of radius r centered at the origin. A path π is *feasible* if the unit disk whose center moves along π does not intersect any obstacle, and if the speed of motion along π is never greater than 1. A *collection* of feasible paths is *feasible* if for any two paths in the collection, the unit disks whose centers move along the paths do not intersect. The objective is to find a feasible collection of maximum cardinality. Let OPT_Ω denote an optimal collection, and let $|OPT_\Omega|$ be the number of paths in OPT_Ω.

Hardness Results. Let $\alpha, \beta > 0$ be positive numbers, arbitrarily small.

THEOREM 3.1. *Unless P=NP, there exists no polynomial-time algorithm to find a feasible collection of $\lceil \alpha |OPT_\Omega| \rceil$ paths for disks of radius β.*

PROOF. *(Sketch.)* Canny and Reif [3] showed that it is NP-hard to establish existence of *one* (thin, radius-0) path between two points, s and t, in a polygonal domain with moving obstacles. The obstacles in [3] do not intersect — a small amount of free space between them exists for paths of relevant homotopy types. After appropriate scaling, there will be just enough space for thickness-β paths (but not for thicker paths). Thus, given an input to Canny and Reif's problem, we may create an input to our problem with the same (but scaled) domain and obstacles, and $\Gamma_I = s$, $\Gamma_O = t$. There exists a (thin) s-t path in the original problem if and only if there exists a β-thick Γ_I-Γ_O path in ours. □

Note that the reduction in [3] uses some *fast* obstacles. We show that our problem is NP-hard even if the obstacles are *static*.

THEOREM 3.2. *The problem of deciding whether $|OPT_\Omega| \geq K$ is NP-hard (weakly NP-hard even for $K = 2$) even if $\mathcal{H}(t) = const$.*

PROOF. *(Sketch.)* Specifying entry and exit times for the paths, together with the speed constraint, implies a restriction on the (two-dimensional, (x, y)) *length* of the paths. Thus, our problem is at least as hard as any problem that asks about existence of K *short* thick paths. The theorem follows from the hardness of establishing existence of K short thick paths amidst static obstacles [15]. □

The above hardness results suggest that in order to get $|OPT_\Omega|$ paths it is not enough just to allow the disks to move faster — one needs also to compromise on the radius. Below we present pseudopolynomial-time algorithm, which finds $|OPT_\Omega|$ paths for disks with radius smaller than 1 and maximum speed greater than 1, under the assumption that the obstacles move with speed at most 1.

Lifting to (x, y, t). We treat Ω as a 3D domain $\Omega = \bigcup_{0 \leq t \leq T} (\Omega(t), t)$ in the (x, y, t)-space. The holes, when moving, sweep a set $\mathcal{X} = \bigcup_{0 \leq t \leq T} (\mathcal{H}(t), t)$ of 3D obstacles. A feasible path π is a curve $(\pi(t), t)$, connecting a point $(\pi(t_I^\pi), t_I^\pi)$ in the rectangle $\Gamma_I \times T_I$ to a point $(\pi(t_O^\pi), t_O^\pi) \in$

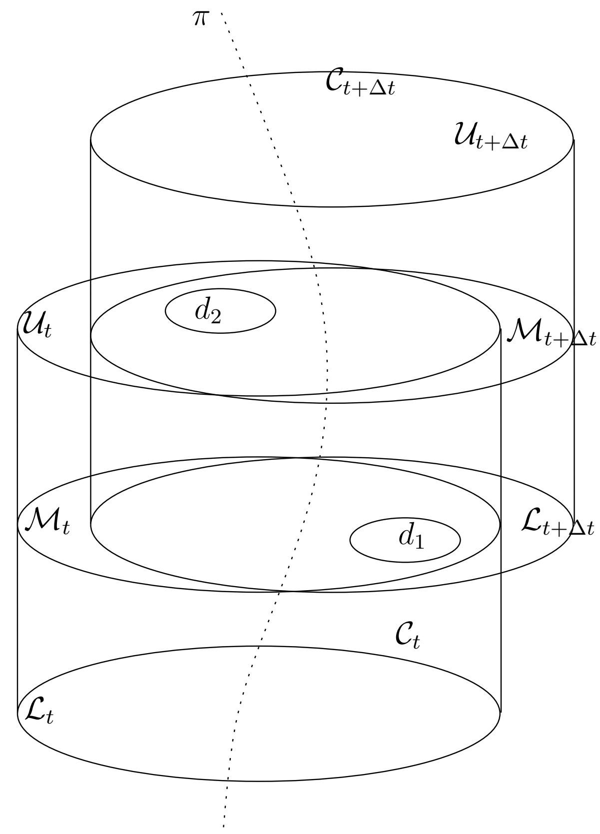

Figure 1: The cylinders inside $< \pi >^{(1)}$.

$\Gamma_O \times T_O$. The Minkowski sum of π with the (two-dimensional, horizontal) unit disk, $< \pi >^{(1)}$, is a "slanted and curved" cylindrical *tube*; for a feasible collection of paths these tubes must not intersect the obstacles or each other.

Obstacles Motion. The condition that the obstacles "move with speed at most 1" formally means that $\mathcal{X}$ has a tangent plane almost everywhere, and that the plane is inclined by at least 45^o to the (x, y)-plane. We assume that the obstacles' motion is such that the following query can be answered in polynomial time: Given time t, two points $a, b \in \Omega$, speed v, and radius r, determine whether a disk of radius r intersects any obstacle when its center starts at point a at time t and moves at constant speed v along a straight line segment to point b. Apart from this, we impose no restrictions on the obstacles' motion, e.g., it is possible that during the motion the obstacles intersect, grow, disappear, etc.

Naming Conventions. All cylinders in this paper have horizontal circular bases. By default, a "cylinder" means a *right* cylinder; oblique cylinders will be called *elementary*, and will be referred to as such. By a *motion* or a *path* of a disk we understand a motion or a path of its center. The *horizontal distance* between two points is the distance between their projections on a horizontal plane. An r-disk is a disk of radius r.

Cylinders Inside a Feasible Path. The following lemma is elementary:

LEMMA 3.3. *Let C be an r-disk; let C' be C, shifted horizontally by $a < 2r$. Then the intersection $C \cap C'$ contains an $(r - a/2)$-disk.*

Let $\Delta t < 2/3$ be some constant. Let π be a feasible path. The next lemma shows that for any $t \in [t_I^\pi, t_O^\pi]$, there exists a "chunk" of obstacle-free space around $\pi(t)$.

LEMMA 3.4. *The cylinder $< \pi(t) >^{(1-\Delta t)} \times [t-\Delta t, t+\Delta t]$ is obstacle-free.*

PROOF. At t, the unit disk centered at $\pi(t)$ is obstacle-free. Since the obstacles' speed is bounded by 1, no obstacle could have been (resp., will be) closer than $1-\Delta t$ to $\pi(t)$ during the Δt-long time interval $[t-\Delta t, t]$ (resp., $[t, t+\Delta t]$). □

We denote the obstacle-free cylinder whose existence is established in the above lemma by $\mathcal{C}_t^\pi$. Let $\mathcal{L}_t^\pi$ (resp., $\mathcal{M}_t^\pi$, $\mathcal{U}_t^\pi$) be the cross-sections of $\mathcal{C}_t^\pi$ by the horizontal plane at $t - \Delta t$ (resp., t, $t + \Delta t$); i.e., $\mathcal{L}_t^\pi$ and $\mathcal{U}_t^\pi$ are the bases of $\mathcal{C}_t^\pi$. Refer to Fig. 1. (We assume that if $t - \Delta t < t_I^\pi$ (resp., $t + \Delta t > t_O^\pi$), the part of $\mathcal{C}_t^\pi$ below (resp., above) $\mathcal{M}_t^\pi$ is chopped off.)

Let $R = \frac{1}{3} - \frac{1}{2}\Delta t$; let $D = \frac{4}{3} - \Delta t$. (This is an important piece of notation; R and D will be used a lot throughout the rest of the paper.)

LEMMA 3.5. *Inside the intersection $\mathcal{M}_t^\pi \cap \mathcal{L}_{t+\Delta t}^\pi$ there exists a $3R$-disk.*

PROOF. Since the speed of motion along a feasible path is at most 1, $\mathcal{L}_{t+\Delta t}^\pi$ is $\mathcal{M}_t^\pi$ shifted horizontally by at most Δt. The lemma follows now from Lemma 3.3. □

Let $d_1 \subset (\mathcal{M}_t^\pi \cap \mathcal{L}_{t+\Delta t}^\pi)$ (resp., $d_2 \subset (\mathcal{M}_{t+\Delta t}^\pi \cap \mathcal{U}_t^\pi)$) be an R-disk, fully lying inside the intersection of the cross-sections (Fig. 1); let c_1 (resp., c_2) be the center of the disk.

LEMMA 3.6. *The horizontal distance between c_1 and c_2 is at most D.*

PROOF. Let d_2' be the projection of d_2 onto the plane of d_1. Since both d_2' and d_1, being R-disks, lie inside the intersection of $(1 - \Delta t)$-disks, the distance between their centers is at most $2(1 - \Delta t) - 2R = D$. □

LEMMA 3.7. *d_1 can be moved to d_2 by a straight-line motion with speed at most $d/\Delta t$ without intersecting any obstacle.*

PROOF. Connect d_1 to d_2 by tube $< c_1c_2 >^{(R)}$ — the convex hull of the disks. Since both disks lie inside the (obstacle-free) cylinder $\mathcal{C}_t^\pi$, the tube does not intersect any obstacle. By Lemma 3.6, the speed of motion along c_1c_2 is at most $D/\Delta t$. □

Slicing the Time. Assume for simplicity that $T/\Delta t$ is an integer, M. Partition the planning horizon, $[0, T]$, into M intervals of length Δt. Let $t_0 \dots t_M$ denote the intervals' endpoints; $t_m = m\Delta t, m \in \{0 \dots M\}$. Applying Lemmas 3.4, 3.5, 3.7, at every interval, we get:

COROLLARY 3.8. *Inside a feasible tube $< \pi >^{(1)}$ there exists a set of cylinders $\{\mathcal{C}_m^\pi\}_{m=1}^M$ with the following properties:*

1: *the height of $\mathcal{C}_m^\pi$ is $2\Delta t$, the radius is $1 - \Delta t$*

2: *$\mathcal{M}_m^\pi$ is in the plane $t = t_m$; hence, $\mathcal{L}_{m+1}^\pi$ and $\mathcal{U}_{m-1}^\pi$ are also in the plane $t = t_m$*

3: *the intersection $\mathcal{M}_m^\pi \cap \mathcal{L}_{m+1}^\pi$ contains a $3R$-disk; hence, so does $\mathcal{M}_{m+1}^\pi \cap \mathcal{U}_m^\pi$*

4: *for any R-disks $d_1 \subset (\mathcal{M}_m^\pi \cap \mathcal{L}_{m+1}^\pi)$ and $d_2 \subset (\mathcal{M}_{m+1}^\pi \cap \mathcal{U}_m^\pi)$ there exists a straight-line motion, with speed at most $D/\Delta t$, that moves d_1 to d_2 without intersecting any obstacle.*

Disk Packings. Let $\mathcal{P}_m$, for $m = 0 \dots M$, be a *maximal packing* of R-disks in $\Omega(t_m)$; let $\mathcal{P} = \cup_m \mathcal{P}_m$. (A maximal packing is a set of disjoint disks such that none of the disks intersects P or a hole in $\mathcal{H}(t_m)$, and no more disks can be placed without violating this.) Suppose that the packings satisfy the following (Γ_I, Γ_O)-*coverage* property: If $t_m \in T_I$ (resp., $t_m \in T_O$), in $\mathcal{P}_m$ a *maximum possible* number of disks is placed with centers on Γ_I (resp., Γ_O). (Following [13], we assume that Ω has Riemann flaps attached along Γ_I, Γ_O so that it is possible to place the *centers* of the disks along Γ_I, Γ_O.) Assume for simplicity that even if T_I and T_O overlap, the disks, placed along Γ_I and Γ_O for $t_m \in T_I \cap T_O$, do not intersect. The next lemma shows how to find a packing with (Γ_I, Γ_O)-coverage, and with an additional property of having at least one disk with center in Γ_I and at least one disk with center in Γ_O inside *any* feasible path in Ω; the property ensures that no path will be "lost" because of the time discretization.

LEMMA 3.9. *Let $\mathcal{P}_{\Gamma_I}$ (resp., $\mathcal{P}_{\Gamma_O}$) be a maximum-cardinality set of disjoint R-disks with centers along Γ_I (resp., Γ_O). Then for every feasible path π in Ω there exists a disk $d_I \in \mathcal{P}_{\Gamma_I}$ (resp., a disk $d_O \in \mathcal{P}_{\Gamma_O}$) such that $d_I \subset < \pi >^{(1)}$ (resp., $d_O \subset < \pi >^{(1)}$).*

PROOF. Let $\pi(t_I^\pi) = s$, i.e., π enters Ω through a point $s \in \Gamma_I$. Since the cylinder $< s >^{(1-\Delta t)} \times [t_I^\pi - \Delta t, t_I^\pi + \Delta t]$ has height $2\Delta t$, it is intersected by the plane $t = t_m$ for some $m = 0 \dots M$. By Lemma 3.4, the cylinder is obstacle-free, and lies fully inside $< \pi >^{(1)}$. Since the radius of the cylinder is $1 - \Delta t > R$, there exists a disk from $\mathcal{P}_{\Gamma_I}$ lying inside the cross-section of the cylinder by the plane $t = t_m$. The existence of d_O is established analogously. It is easy to see that a maximum packing of disks can be achieved by greedily packing the disks, starting from an endpoint of Γ_I, Γ_O. □

Motion Graph. Let $G = (\mathcal{P}, E)$ be the *motion graph* — a directed graph with vertices being the disks in the packings, and edges defined as follows. For disks $d_1, d_2 \in \mathcal{P}$ there is an edge $(d_1, d_2) \in E$ whenever the vertical distance between d_1 and d_2 is Δt (i.e., $d_1 \in \mathcal{P}_m$, $d_2 \in \mathcal{P}_{m+1}$ for some m), *and* there exists a straight-line motion, with speed at most $\frac{4}{3\Delta t} - 1$, that moves d_1 to d_2 without intersecting any obstacle (Fig. 2). Add a super-source vertex S (resp., super-sink vertex T) to $\mathcal{P}$; connect S (resp., T) to the disks placed along Γ_I (resp., Γ_O) for $t_m \in T_I$ (resp., $t_m \in T_O$). By a *path in G* we will always mean an S-T path.

We will assume that G is embedded in (x, y, t) space with vertices at the disks centers. Depending on the context, by an *edge e* of G we will mean four related things: (1) the directed edge of the graph; (2) the vector, directed upward, in the (x, y, t)-space; (3) the (undirected) segment in the (x, y, t)-space; and, (4) the oblique cylinder $< e >^{(R)}$, which

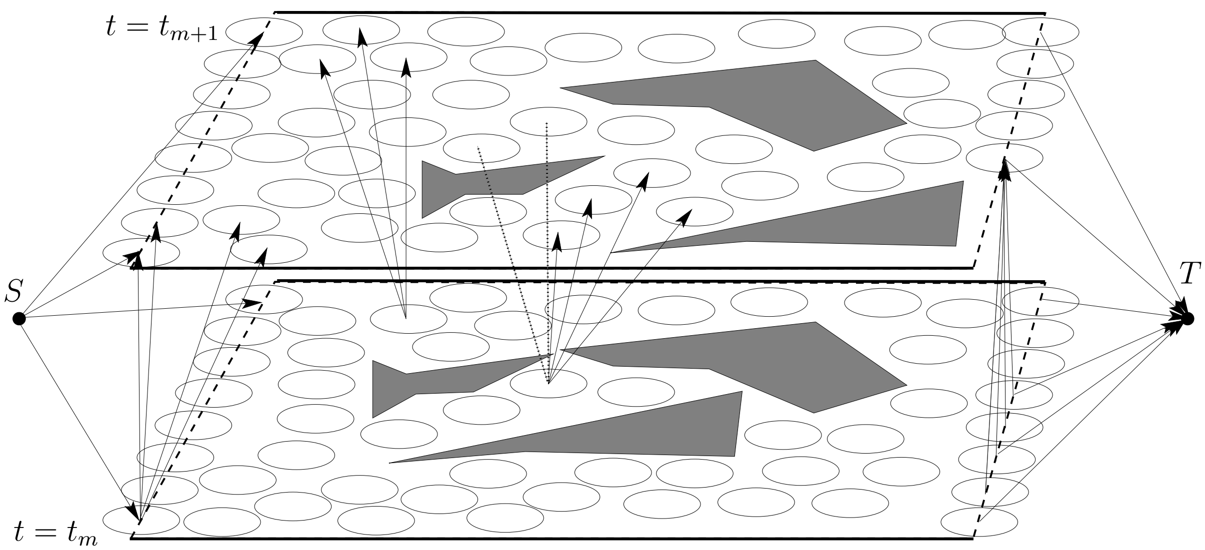

Figure 2: P is a rectangle; Γ_I, Γ_O are dashed. The edges of G, adjacent to a disk in the packing, connect the disk to all disks that it can reach within time Δt without intersecting obstacles. Some of the edges are shown. The dotted segments are not edges because their elementary cylinders are intersected by an obstacle path.

we call an *elementary cylinder*. The specific meaning will be clear from the context; e.g., in the expression $< e >^{(r)}$ for $r \neq R$, the edge e will necessarily mean a segment (meaning (3)). A path in G will be identified with the tube, obtained by inflating the edges of the path (excluding the edges adjacent to S and T).

By construction, any path π^G in G is an obstacle-free path for an R-disk; $< \pi^G >^{(R)} \cap \mathcal{X} = \emptyset$. In the next lemma we show that, conversely, there also exists a path in G for each feasible thickness-1 path in Ω.

LEMMA 3.10. *Let π be a feasible path in Ω. There exists a path π^G in G such that $< \pi^G >^{(R)} \subset < \pi >^{(1)}$.*

PROOF. Since $\mathcal{P}_m$ is maximal, any obstacle-free $3R$-disc in $\Omega(t_m)$ contains at least one disc from $\mathcal{P}_m$. In particular, there exist a disk d_1 from $\mathcal{P}_m$ inside $\mathcal{M}^\pi_m \cap \mathcal{L}^\pi_{m+1}$ and a disk $d_2 \in \mathcal{P}_{m+1}$ inside $\mathcal{M}^\pi_{m+1} \cap \mathcal{U}^\pi_m$ (Corollary 3.8, property 3). By Corollary 3.8, property 4, $(d_1, d_2) \in E$. By the (Γ_I, Γ_O)-coverage property, there exists an edge, connecting some disk lying inside $< \pi >^{(1)}$ to S, and an edge connecting some disk to T. □

In what follows we attempt to reduce computing $|OPT_\Omega|$ paths to a version of maxflow in a modified graph G. Although we do not succeed straightforwardly, we show how to find $|OPT_\Omega|$ paths by compromising on the disks' radii and maximum speed.

MaxFlow in G — a Naïve Attempt. In the sequel whenever we speak of a *flow*, we will mean an *integral S-T* flow in G. Assign capacity 1 to every vertex and every edge of G. Let $\mathbf{f}(G)$ denote the maximum flow; let $|\mathbf{f}(G)|$ be its value. By the Flow Decomposition Theorem, $\mathbf{f}(G)$ can be decomposed into a set of $|\mathbf{f}(G)|$ vertex-disjoint paths. By Lemma 3.10, $|\mathbf{f}(G)| \geq |OPT_\Omega|$. Although no path in the set $\mathbf{f}(G)$ intersects an obstacle, nothing prevents the paths in $\mathbf{f}(G)$ from intersecting each other. Hence, just finding a maxflow in G is not enough to get a feasible collection of $|OPT_\Omega|$ paths: the paths must be prohibited from using intersecting elementary cylinders. Next we show how the flow problem can be modified in order to avoid such intersections.

Flow with Forbidden Pairs — Less Naïve, but Hopelessly Hard. Let $\mathcal{F}$ be the set of pairs of intersecting elementary cylinders; for $e_1, e_2 \in E$, the pair (e_1, e_2) is in $\mathcal{F}$ iff $< e_1 >^{(R)} \cap < e_2 >^{(R)} \neq \emptyset$. We say that a flow *respects* $\mathcal{F}$ if for any pair of edges from $\mathcal{F}$ at most one edge from the pair is used (hence—by integrality—saturated) by the flow. The *MaxFlow with Forbidden Pairs* (MFFP) problem is that of finding a maximum flow respecting a set of pairs of edges. Let $\text{MFFP}(G, \mathcal{F})$ be the maximum flow that respects $\mathcal{F}$, and let $|\text{MFFP}(G, \mathcal{F})|$ be its value.

LEMMA 3.11. $|\text{MFFP}(G, \mathcal{F})| \geq |OPT_\Omega|$.

PROOF. By Lemma 3.10, there exists a path in G per tube in OPT_Ω. Since each path stays inside its own tube, the paths yield a feasible collection of $|OPT_\Omega|$ thickness-R paths. Since the collection is feasible, no pair of edges from $\mathcal{F}$ is simultaneously used by them. □

Thus, to get a collection of $|OPT_\Omega|$ paths it would be enough to solve MFFP. Unfortunately, as the next proposition shows, even *approximating* MFFP is hard. This means that obtaining $|OPT_\Omega|$ paths is impossible without further compromise (we thank Jim Orlin for this insight).

PROPOSITION 3.12. *Unless P=NP* MFFP *is not approximable to within factor $\Omega(\sqrt{\text{degree of } S})$ even for "layered" graphs with at most one forbidden pair per layer.*

PROOF. *(Sketch.)* Given an instance of MAX INDEPENDENT SET [4] on a graph G' with n' nodes, create n' disjoint S-T paths, each of length $\binom{n'}{2}$; a vertex of G' corresponds to a path. Create also a forbidden pair per edge of G'. Now maxflow value in the created graph = size of the independent set in G'. The proposition follows from the hardness of approximating MAX INDEPENDENT SET [4]. □

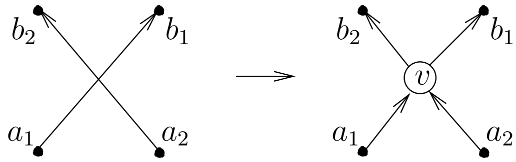

Figure 3: The edges e_1, e_2 from a forbidden pair are replaced by four edges: $(a_1, v), (a_2, v), (v, b_1), (v, b_2)$. This preserves all connections in G, but also introduces connections $a_1 \to b_2$ and $a_2 \to b_2$, which may not have been present in G.

Modified G. An upper bound on MFFP$(G, \mathcal{F})$ can be obtained as follows. For every pair $(e_1, e_2) \in \mathcal{F}$ create a vertex $v = v(e_1, e_2)$, and replace the edges e_1, e_2 by (a_1, v), (a_2, v), (v, b_1), and (v, b_2), where $(a_1, b_1) = e_1, (a_2, b_2) = e_2$ (Fig. 3); assign capacity 1 to v and to the new edges. Let $G_{\mathcal{F}}$ denote G after such replacement has been made for all pairs in $\mathcal{F}$.

Lemma 3.13. $|\mathbf{f}(G_{\mathcal{F}})| \geq |\mathbf{f}(G)| \geq |\text{MFFP}(G, \mathcal{F})|$.

Proof. Since $G_{\mathcal{F}}$ preserves all connections between vertices of G, for any flow in G there exists a flow in $G_{\mathcal{F}}$; thus, $|\mathbf{f}(G_{\mathcal{F}})| \geq |\mathbf{f}(G)|$. Of course, an $\mathcal{F}$-respecting flow is a flow; thus, $|\mathbf{f}(G)| \geq |\text{MFFP}(G, \mathcal{F})|$. □

Unfortunately, $G_{\mathcal{F}}$ has some new connections, not present in G. Specifically, a flow in $G_{\mathcal{F}}$ may go via a_1-v-b_2 or a_2-v-b_1, even when that was not possible in G, i.e., if (a_1, b_2) or (a_2, b_1) were not in E. (Thus, in general $|\text{MFFP}(G_{\mathcal{F}}, \mathcal{F})| \neq |\text{MFFP}(G, \mathcal{F})|$, which is not surprising in view of Proposition 3.12.) Next we show that at the expense of further compromising on the disks' radius and maximum speed, we actually may allow such connections to exist inside pairs of "deeply penetrating" elementary cylinders, without tampering with paths following cylinders that overlap only slightly.

Mutual Penetration of Elementary Cylinders. We show that since elementary cylinders are "squashed" (i.e., their radius is large in comparison with their height and the horizontal distance between their bases' centers), if two of them deeply penetrate, then the penetration occurs far from their bases, and no other cylinder deeply penetrates either of the two. This implies that a disk of "small" radius can be moved with "moderate" speed from the lower base of one cylinder to the upper base of *the other.*

Let $e_1 = (a_1, b_1)$, $e_2 = (a_2, b_2)$, $e_1, e_2 \in E$ be two elementary cylinders from one layer of G (recall that we view edges of G as being "thick" cylinders). Let p_i be a point moving along e_i, for $i = 1, 2$; let $p_i(\tau)$ be the position of p_i when it is at distance τ from the plane containing the lower bases of e_1, e_2. Let $0 < \delta < 2R^2$ be an arbitrary constant. (Remark: δ is actually a *distance*; the 2 in "$2R^2$" has units of distance^{-1} and comes from 2/(radius of the unit ball)).

Definition 3.14. *The* penetration *of e_1, e_2, denoted $\mathsf{P}(e_1, e_2)$, is the minimum distance between the cross-sections of the axes of e_1, e_2 by a horizontal plane; i.e., $\mathsf{P}(e_1, e_2) = \min_\tau |p_1(\tau)p_2(\tau)|$. The* height *of the penetration, denoted $\mathsf{h}(e_1, e_2)$, is the value of τ at which the minimum is attained. We say that e_1, e_2* deeply penetrate *if $\mathsf{P}(e_1, e_2) \leq \delta$.*

Note that according to our definition, if e_1, e_2 overlap a lot, $P(e_1, e_2)$ is *small*; this is somewhat counterintuitive, but makes the following formulae simpler.

Let $\mathbf{l} = a_2 - a_1, \mathbf{u} = b_2 - b_1$; let $l = |\mathbf{l}|$, $u = |\mathbf{u}|$, (Fig. 4, left). Let $\widehat{\mathbf{vw}}$ be the angle between vectors $\mathbf{v}$ and $\mathbf{w}$.

Lemma 3.15. *If e_1, e_2 deeply penetrate, $\widehat{\mathbf{lu}} \in [150^o, 210^o]$.*

Proof. Suppose, w.l.o.g., that $l \geq u$. Since $p_2(\tau) - p_1(\tau) = \mathbf{l} + \frac{\tau}{\Delta t}(\mathbf{u} - \mathbf{l})$, we have

$$h(e_1, e_2) = \frac{\mathbf{l} \cdot (\mathbf{l} - \mathbf{u})}{|\mathbf{l} - \mathbf{u}|^2} \Delta t \quad , \quad P(e_1, e_2) = \frac{|\mathbf{l} \times \mathbf{u}|}{|\mathbf{l} - \mathbf{u}|} \qquad .$$

Since $|e_1|, |e_2| \leq 1$, we have $|e_1 - e_2| \leq 2$ (here, e_1, e_2 are treated as vectors in (x, y, t)-space). Since e_1, e_2 deeply penetrate, and $|\mathbf{l} - \mathbf{u}| = |e_1 - e_2|$, we have $|\sin(\widehat{\mathbf{lu}})| \leq 2\delta/(lu)$; i.e.,

$$\begin{aligned} \widehat{\mathbf{lu}} \quad \in \quad & [0, \arcsin \tfrac{2\delta}{lu}] \\ \cup\, & [180^o - \arcsin \tfrac{2\delta}{lu}, 180^o + \arcsin \tfrac{2\delta}{lu}] \\ \cup\, & [360^o - \arcsin \tfrac{2\delta}{lu}, 360^o]. \end{aligned} \tag{1}$$

On the other hand, since the bases of e_1, e_2 are disjoint, to have $P(e_1, e_2) \leq \delta < 2R^2 < 2R$, it is necessary that $0 < h(e_1, e_2) < \Delta t$, i.e., that $0 < \mathbf{l} \cdot (\mathbf{l} - \mathbf{u}) < (\mathbf{l} - \mathbf{u})^2$, which implies $\cos(\widehat{\mathbf{lu}}) \leq u/l$, or $\widehat{\mathbf{lu}} \in [90^o - \arccos \frac{u}{l}, 270^o + \arccos \frac{u}{l}]$, which is consistent only with the middle set in (1) for $\delta < u^2/2$. Since the bases of e_1, e_2 are disjoint, $u \geq 2R$; hence, $\delta < 2R^2$ implies $\delta < u^2/2$. Thus, $|\widehat{\mathbf{lu}} - 180^o| < \arcsin \frac{2\delta}{lu}$, which is less than 30^o since $l, u \geq 2R$. □

Lemma 3.16. *If $\mathsf{P}(e_1, e_2) \leq \delta$, then $\forall e \neq e_1, e_2$, we have $\mathsf{P}(e, e_1) \geq \frac{3}{4}\delta$.*

Proof. *(Sketch.)* Suppose otherwise. Let $e = (a, b)$. Let $\mathbf{l}' = a - a_1$, $\mathbf{u}' = b - b_1$; let $l' = |\mathbf{l}'|$, $u' = |\mathbf{u}'|$ (see Fig. 4, left). The penetration implies that $l, l' < 3R$. Thus, $\widehat{\mathbf{ll}'}$ is minimum when a is in one of the extremal positions in Fig. 4, center; in both of them, $\widehat{\mathbf{ll}'} \geq \arccos(3/4) > 40^o$. Similarly, $\widehat{\mathbf{uu}'} > 40^o$. Together with $\widehat{\mathbf{lu}}, \widehat{\mathbf{l}'\mathbf{u}'} \in [150^o, 210^o]$ and $l, u, l', u', |\mathbf{l} - \mathbf{u}|, |\mathbf{l}' - \mathbf{u}'| \geq 2R$ this implies (by case analysis of how $\mathbf{l}, \mathbf{l}', \mathbf{u}, \mathbf{u}'$ are situated) that $|\mathbf{l}' \times \mathbf{u}'| \geq \frac{3}{4}\delta|\mathbf{l}' - \mathbf{u}'|$. □

Lemma 3.17. *If e_1, e_2 deeply penetrate,* $(2R - \delta)/(2D/\Delta t) \leq \mathsf{h}(e_1, e_2) \leq \Delta t - (2R - \delta)/(2D/\Delta t)$.

Proof. Since the lower bases of e_1, e_2 are disjoint, when p_1, p_2 start moving along e_1, e_2, the distance between them is at least $2R$. Since the slope of e_1, e_2 is at most $D/\Delta t$, p_1, p_2 move towards each other not faster than $2D/\Delta t$. Similarly, the penetration must end before reaching the upper bases of e_1, e_2. □

The next lemma provides the basis for our "solution" to MFFP, given below.

Lemma 3.18. *If e_1, e_2 deeply penetrate, a disk of radius $\delta/8$, moving at speed at most $10/\Delta t$, may be moved from a_1 to b_2 without intersecting the cylinder $< e >^{(\delta/8)}$ for any $e \in E$, $e \neq e_1, e_2$.*

Proof. Let $p_1 = p_1(h(e_1, e_2))$, $p_2 = p_2(h(e_1, e_2))$; let c be the midpoint of p_1p_2. Let p be the intersection of (the axis of) e with the plane $t = h(e_1, e_2)$. By Lemma 3.16, $|p_1p| \geq \frac{3}{4}\delta$. Since $|p_1c| < \delta/2$, we have $|cp| \geq \delta/4$, which

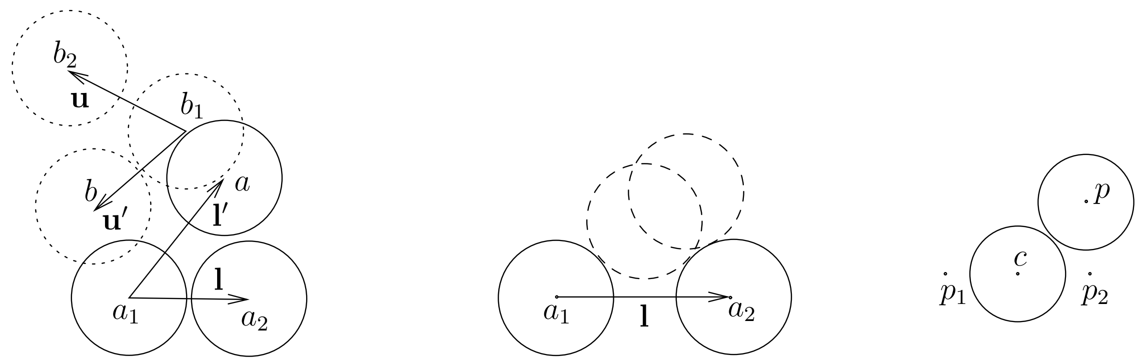

Figure 4: Left: If $\mathsf{P}(e_1, e_2) < \delta$, then l and u "point in different directions". For clarity, the figure is not drawn to scale. In reality, $|a_2 b_2|$ must be much smaller in comparison with the radii; if drawn to scale, there must be much more overlap between the disks, and the figure becomes too cluttered. Center: The extremal positions of a, leading to the maximum possible value of $\widehat{\mathbf{l}\mathbf{l}'}$. Right: Cross-section with the plane $t = \mathsf{h}(e_1, e_2)$. $p_i = p_i(\mathsf{h}(e_1, e_2))$. $|p_1 p| \geq \frac{3}{4}\delta, |p_1 c| < \delta/2 \Longrightarrow |cp| \geq \delta/4 \Longrightarrow < p_1 >^{(\delta/8)} \cap < e >^{(\delta/8)} = \emptyset$.

means that a $\delta/8$-disk, centered at c, does not intersect $< e >^{(\delta/8)}$ (Fig. 4, right). Also by Lemma 3.16, a $\delta/8$-disk, centered at a_1, does not intersect $< e >^{(\delta/8)}$ either. Thus, the tube $< a_1 c >^{(\delta/8)}$—the convex hull of the disks—does not intersect $< e >^{(\delta/8)}$.

Since c is inside e_1, the horizontal distance from a_1 to c is at most $|a_1 b_1| \leq D$ (Lemma 3.6). By Lemma 3.17, a disk, starting at a_1, has time at least $(2R - \delta)/(\frac{2D}{\Delta t})$ to get to c; thus the disk may move with speed at most $\frac{2}{\Delta t}\frac{D^2}{2R-\delta} < 10/\Delta t$ for $\Delta t < 1/2$. Similarly, the disk may be moved from c to b_2. $\square$

We will call the operation of moving the disk as described in the above lemma *tube bending*. It is easy to see that tubes, bent inside different pairs of deeply penetrating cylinders, do not intersect.

A "Solution" to MFFP— **Incorrect, but Fixable via Bending.** Let $\mathcal{F}^* \equiv \{e_1, e_2 \in E \,|\, P(e_1, e_2) < \delta\}$ be the set of forbidden pairs of deeply penetrating elementary cylinders. Decompose $\mathbf{f}(G_{\mathcal{F}^*})$ into a set $\mathcal{K}$ of $|\mathbf{f}(G_{\mathcal{F}^*})|$ paths. (Here, similar to $G_{\mathcal{F}}$, we denote by $G_{\mathcal{F}^*}$ the graph in which there is a supernode for every pair of forbidden edges in $\mathcal{F}^*$, and by $\mathbf{f}(G_{\mathcal{F}^*})$ the maximum flow in it.) We say that a path $\pi \in \mathcal{K}$ *wiggles in* $\mathbf{f}(G_{\mathcal{F}^*})$ if it has a subpath a_1-v-b_2, where $e_1 = (a_1, b_1), e_2 = (a_2, b_2)$, $(e_1, e_2) \in \mathcal{F}^*$, $v = v(e_1, e_2)$ is the supernode for (e_1, e_2), but (a_1, b_2) is *not* in E. The paths that do not wiggle in $\mathbf{f}(G_{\mathcal{F}^*})$ correspond to a feasible collection of thickness-$\delta/8$ tubes in Ω; indeed, two tubes could intersect only if they followed a forbidden pair, which is prohibited by the capacity of the supernode. On the other hand, $\delta/8$-thick tubes, following wiggling paths, may be bent as in Lemma 3.18 without conflicting with each other and the tubes that correspond to paths, not wiggling in $G_{\mathcal{F}^*}$. Thus,

LEMMA 3.19. *A feasible collection of* $|\mathbf{f}(G_{\mathcal{F}^*})|$ *thickness-*$\delta/8$ *paths may be built from* $\mathbf{f}(G_{\mathcal{F}^*})$.

Since fewer pairs conflict in $G_{\mathcal{F}^*}$ than in $G_{\mathcal{F}}$ (due to $\mathcal{F}^* \subseteq \mathcal{F}$), we have $|\mathbf{f}(G_{\mathcal{F}^*})| \geq |\mathbf{f}(G_{\mathcal{F}})|$. Together with Lemmas 3.13 and 3.11, this implies that $|\mathbf{f}(G_{\mathcal{F}^*})| \geq |OPT_{\Omega}|$.

Putting Things Together. The algorithm MAXTHICK-PATHS in Fig. 5 implements the steps described above with δ chosen to be $2(\frac{1}{3} - \frac{1}{2}\Delta t)^2$ — the maximum allowed value. Our main result is:

THEOREM 3.20. *For any $\Delta t < 1/2$, algorithm* MAXTHICK-PATHS *computes a collection of (at least) $|OPT_{\Omega}|$ feasible paths for radius-$(\frac{1}{3} - \frac{1}{2}\Delta t)^2$ disks, moving with speed at most $10/\Delta t$. The algorithm runs in time polynomial in $T/\Delta t$ and N, where N is the value of largest coordinate of Ω.*

Extensions

Our approach can be extended to higher dimensions and to other shapes of the moving objects (as long as the motion is translational). It also allows us to incorporate additional constraints on the paths. For instance, in ATM, it may be undesirable to route a huge number of aircraft into the airspace at the same instant of time. To address this, instead of connecting S directly to Γ_I-nodes at all time levels, insert an arc and a node between S and the Γ_I-nodes at every level; the capacity of the arc will bound the number of paths that enter Ω at any given time. In addition, it is trivial to impose monotonicity on the paths: just keep the relevant arcs in G. We may also allow the obstacles to move faster. Their maximum speed affects the approximation guarantees of our algorithm in a quantifiable way. Finally, we may allow there to be multiple sources/sinks, both on the boundary and inside Ω.

Acknowledgements

We are indebted to an anonymous reviewer for many insightful comments and corrections. E. Arkin is partially supported by NSF (CCF-0431030, CCF-0729019). J. Mitchell is partially supported by NSF (CCF-0431030, CCF-0528209, CCF-0729019), NASA Ames, and Metron Aviation.

4. REFERENCES

[1] N. Aggarwal and K. Fujimura. Motion planning amidst planar moving obstacles. In *Proc. International Conference on Robotics and Automation*, pages 2153–2158, 1994.

[2] J. Canny, A. Rege, and J. Reif. An exact algorithm for kinodynamic planning in the plane. *Discrete Computational Geometry*, 6:461–484, 1991.

Algorithm MaxThickPaths $(\Omega(t), T_I, T_O)$
Input. Domain Ω with moving obstacles; entry and exit intervals T_I, T_O; user-defined parameter $\Delta t < 1/2$.
Output. A feasible collection of $|OPT_\Omega|$ paths for radius-$(\frac{1}{3} - \frac{1}{2}\Delta t)^2$ disks moving with maximum speed $10/\Delta t$.

$T \leftarrow \max\{T_O\} - \min\{T_I\}$ $\quad \triangleright$ Planning horizon.
$M \leftarrow T/\Delta t$ $\quad \triangleright$ It is assumed for simplicity that M is an integer.
$R \leftarrow \frac{1}{3} - \frac{1}{2}\Delta t$
for $m = 0$ **to** M
$t_m \leftarrow m \cdot \Delta t$
$\mathcal{P}_m \leftarrow$ maximal packing of radius-R disks in $\Omega(t_m)$ $\quad \triangleright$ with (Γ_I, Γ_O)-coverage.
endfor
$\mathcal{P} \leftarrow \bigcup_m \mathcal{P}_m$
$G \leftarrow$ motion graph on $\mathcal{P}$
$\mathcal{F}^* \leftarrow \{(e_1, e_2) \in E \,|\, \mathsf{P}(e_1, e_2) < 2R^2\}$ $\quad \triangleright$ Forbidden pairs of deeply penetrating elementary cylinders.
$G_{\mathcal{F}^*} \leftarrow G$ with supernodes for pairs in $\mathcal{F}^*$
$\mathbf{f}(G_{\mathcal{F}^*}) \leftarrow$ maximum number of paths in $G_{\mathcal{F}^*}$
Bend the tubes whose paths wiggle in $\mathbf{f}(G_{\mathcal{F}^*})$

Figure 5: Algorithm MaxThickPaths.

[3] J. Canny and J. H. Reif. New lower bound techniques for robot motion planning problems. In *Proc. 28th IEEE Symposium on Foundations of Computer Science*, pages 49–60, 1987.

[4] P. Crescenzi and V. Kann. A compendium of NP optimization problems. In *Complexity and Approximation. Combinatorial Optimization Problems and Their Approximability Properties*. Springer, Berlin, 1999.

[5] K. Fujimura. Motion planning amid transient obstacles. *The International Journal of Robotics Research*, 13(5):395–407, 1994.

[6] K. Fujimura. Time-minimal paths amidst moving obstacles in three dimensions. *Theoretical Computer Science*, 270(1-2):421–440, 2002.

[7] K. Fujimura and H. Samet. Planning a time-minimal motion among moving obstacles. *Algorithmica*, 10(1):41–63, 1993.

[8] J. Krozel, J. S. B. Mitchell, V. Polishchuk, and J. Prete. Capacity estimation for level flight with convective weather constraints. *Air Traffic Control Quarterly*, 15(3):209–238, 2007.

[9] J. Krozel, J. S. B. Mitchell, V. Polishchuk, and J. Prete. Airspace capacity estimation with convective weather constraints. In *Proc. AIAA Guidance, Navigation, and Control Conference*, Aug 2007.

[10] J. Krozel, S. Penny, J. Prete, and J. S. B. Mitchell. Comparison of algorithms for synthesizing weather avoidance routes in transition airspace. In *Proc. AIAA Guidance, Navigation and Control Conference*, Aug 2004.

[11] J.-C. Latombe. *Robot Motion Planning*. Kluwer, Boston, 1991.

[12] Z. Li and J. F. Canny, editors. *Nonholonomic Motion Planning*. Kluwer, Norwell, MA, 1992.

[13] J. S. B. Mitchell. On maximum flows in polyhedral domains. *Journal of Computer and System Sciences*, 40:88–123, 1990.

[14] J. S. B. Mitchell. Shortest paths and networks. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 24, pages 445–466. CRC Press LLC, 1997.

[15] J. S. B. Mitchell and V. Polishchuk. Thick non-crossing paths and minimum-cost flows in polygonal domains. In *Proc. 23rd ACM Symposium on Computational Geometry*, pages 56–65, 2007.

[16] J. S. B. Mitchell, V. Polishchuk, and J. Krozel. Airspace throughput analysis considering stochastic weather. In *Proc. AIAA Guidance, Navigation, and Control Conference*, Aug 2006.

[17] D. Nieuwenhuisen, J. P. van den Berg, and M. H. Overmars. Efficient path planning in changing environments. In *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems - IROS'07*, 2007.

[18] V. Polishchuk. *Thick Non-Crossing Paths and Minimum-Cost Continuous Flows in Geometric Domains*. PhD thesis, Stony Brook University, Aug 2007.

[19] J. Prete. *Aircraft Routing in the Presence of Hazardous Weather*. PhD thesis, Stony Brook University, Aug 2007.

[20] J. Prete and J. S. B. Mitchell. Safe routing of multiple aircraft flows in the presence of time-varying weather data. In *Proc. AIAA Guidance, Navigation and Control Conference*, Aug 2004.

[21] J. Reif and M. Sharir. Motion planning in the presence of moving obstacles. *Journal of the ACM*, 41(4):764–790, 1994.

[22] G. Strang. Maximal flow through a domain. *Mathematical Programming*, 26:123–143, 1983.

[23] J. van den Berg. *Path Planning in Dynamic Environments*. PhD thesis, Utrecht University, 2007.

[24] J. P. van den Berg, D. Ferguson, and J. Kuffner. Anytime path planning and replanning in dynamic environments. In *Proc. International Conference on Robotics and Automation*, pages 2366–2371, 2006.

[25] J. P. van den Berg, D. Nieuwenhuisen, L. Jaillet, and M. H. Overmars. Creating robust roadmaps for motion planning in changing environments. In *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2415–2421, 2005.

[26] J. P. van den Berg and M. H. Overmars. Prioritized

motion planning for multiple robots. In *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2217–2222, 2005.

[27] J. P. van den Berg and M. H. Overmars. Roadmap-based motion planning in dynamic environments. *IEEE Transactions on Robotics*, 21(5):885–897, 2005.

[28] J. P. van den Berg and M. H. Overmars. Planning the shortest safe path amidst unpredictably moving obstacles. In *Proc. Workshop on Algorithmic Foundations of Robotics*, 2006.

[29] J. P. van den Berg and M. H. Overmars. Kinodynamic motion planning on roadmaps in dynamic environments. In *Proc. IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007.