

Deformable Volumes in Path Planning Applications

Elliot Anshelevich Scott Owens Florent Lamiroux Lydia E. Kavraki
Department of Computer Science, Rice University
Houston, TX 77005

Abstract

This paper addresses the problem of path planning for a class of deformable volumes under fairly general manipulation constraints. The underlying geometric model for the volume is provided by a mass-spring representation. It is augmented by a realistic mechanical model. The latter permits the computation of the shape of the considered object with respect to the grasping constraints by minimizing the energy function of the deformation of the object. Previous research in planning for deformable objects considered the case of elastic plates and proposed a randomized framework for planning paths for plates under manipulation constraints. The present paper modifies and extends the previously proposed framework to handle simple volumes. Our planner builds a roadmap in the configuration space. The nodes of the roadmap are equilibrium configurations of the considered volume under the manipulation constraints, while its edges correspond to quasi-static equilibrium paths. Paths are found by searching the roadmap. We present experimental results that illustrate our approach. Our work finds important applications in industrial settings, in graphics animation, in virtual prototyping, and in medical applications.

1 Introduction

The Problem The problem considered in this paper is an extension of the traditional path planning problem: given a deformable object/robot with known physical characteristics (e.g., an elastic object) and a set of ways that the object can be manipulated (called manipulation constraints hereafter), find a quasi-static path for the object from an initial to a goal configuration. The deformations of the object are prescribed by the manipulation constraints in combination with a physical model of the object. Importantly, in our work we consider the geometry and the physical properties of the manipulated object simultaneously. Our goal is to compute paths that are as close as possible to the paths that we expect to observe in practice. That is why our collision-free paths consist of low-energy equilibrium deformations at all times. In this paper we consider simple volumes, such as pipes and cables, that are manipulated by constraining the position and the orientation of two opposite faces. An




Figure 1: An elastic cable fixed to a base is manipulated around a frame and finally attached to it.

illustration is offered in Figure 1, where a cable attached to a base is manipulated by its free end and is finally attached to a frame. This example could have been extracted from an assembly maintainability study: determining whether the motion of the cable is feasible and showing the possible motions to the engineer could help in the overall design of the assembly. Our work in this paper complements and extends our previous work on planning for elastic plates [9, 10].

Motivation and Related Work Several important applications motivate our research: in industrial settings there is a need to manipulate sheets of metal, pipes that can bend, and cables. In assembly maintainability studies done with virtual prototyping, planning is used to evaluate the design when only the CAD models are available [2]. As many assemblies are designed to be compact and include deformable parts, our studies can be used for design verification. In medical and surgical procedures, flexible catheters are inserted in human vessels [16]. Accurate planning studies may help in choosing the size and properties of the catheter used. An interesting application of our work is in computer generated animation, where our methods can be used to produce the motion of objects with certain physical properties with minimal user input. In that case manipulation constraints can be specified to induce a realistic behavior of the animated object. Additional applications of our work include computer-assisted pharmaceutical drug design, where the internal energy of molecules is considered when computing docking paths [14].

Recent work on the path planning problem has produced several practical planners for robots that consist of rigid parts (for recent articles see [5]). These methods routinely take into account geometric constraints such as joint limits and obstacles. With the exception of the areas of dynamics and control that have guided the design of modern robots, there are few cases where physical constraints and planning have been tightly coupled. The issue of flexibility has been primarily investigated by building and studying flexible robots. Those robots can perform tasks such as hammering a peg into a hole or inserting one end of a wire into a hole while holding the other end. But research in this field deals mainly with the control of these robots and not motion planning (see [11]). Although planning with deformable parts has not been addressed in robotics, there is a large amount of work concerning deformable objects in mechanics [17] where elasticity is a well understood issue. Additionally, graphics applications use deformable models [1, 6, 12, 15]. In subsequent sections, we discuss the models that we borrow from these domains.

Our Approach In our work we combine the geometry and the mechanics of our objects/robots under a unified randomized framework. This framework was introduced in [9, 10] for elastic plates. In previous work deformations were essentially two-dimensional. In the present paper we consider true three dimensional deformations of three dimensional volumes. We model our volumes with a mass-spring representation. Mass spring-models have been used in computer graphics [4] and physically-based modeling (e.g. [7, 13]), but have not been applied to motion planning problems. Our manipulation constraints restrict the position and orientation of two opposite faces. Our planner is a variation of probabilistic roadmaps [8]. At a preprocessing stage, it creates random deformations of the robot by varying the manipulation constraints. Random configurations having the above deformations are then generated and interconnected in a roadmap with low-energy quasi-static paths. Given an initial and a final configuration, the planner connects these to the roadmap and searches the roadmap for a path.

This paper is organized as follows. Section 2 describes the deformation of volumes and our model for representing them. Section 3 explains the path planning algorithm. Section 4 presents the case of an elastic pipe and our experimental results.

2 Volume Deformation

We represent a volume as a lattice of point masses connected by springs. The volume deforms according to specified types of manipulation, such as grasping by robotic arms. We attempt to find physically realistic deformations by minimizing the elastic energy of the volume while observing the given manipulation constraints. The use of a

mass spring model provides a geometric representation for the shape as well as a simple framework for calculating deformation energies. We chose this model, among many other that are possible [1, 3, 7, 18], because one can easily approximate the energy of the object from the continuous mechanical model and because of the simplicity of energy and gradient calculations.

2.1 Continuous Mechanical Model

We use the linear elastic physical model [17] as the basis for our spring-mass energy model. For the physical model, we must first choose some embedding of our volume V into \mathbb{R}^3 . We call the undeformed shape U . This is the shape that V assumes when nothing affects it. For example, if V is an elastic cube, the undeformed shape is a cube. Let $\gamma : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ be a diffeomorphism (that is a smooth bijection whose inverse is also smooth) which represents a deformation of our volume from its undeformed state U . The deformed volume is the image $\gamma(U)$. The linear elastic model associates a certain amount of energy with each deformation. For a given deformation γ , define the energy density ψ at a point $v \in \gamma(U)$ by

$$\psi(v) = \frac{1}{2}\lambda(\text{trace}(e))^2 + \mu \text{trace}(e^2), \quad (1)$$

where trace of a matrix (a_{ij}) is $\sum_{i=0}^n a_{ii}$ and

$$e = \frac{1}{2}(F^T F - I) \quad (2)$$

is the Green-Lagrange tensor of the deformation. Here F is the matrix of partial derivatives of γ evaluated at $\gamma^{-1}(v)$ and I is the identity matrix. λ and μ are constants determined by the physical properties of the material [17].

The energy of γ is then

$$\int_{\gamma(U)} \psi(t) dt. \quad (3)$$

2.2 Discrete Spring Model

The space of all possible deformations on a volume is infinite dimensional, but heavily constrained by the requirement to consider minimal energy deformations. We approximate it with a discretization using the mass spring model (Figure 2). We require that the total energy of the springs is an approximation of Eq. (3).

The spring model divides the volume into a set of boxes B with lattice points as the vertices. Noticing that there is one box per point mass (ignoring the boundary special cases), we associate each box with its bottom, left, back point p . We associate six springs with each box, three *straight* springs connecting p to the masses above, right and in front of it and three *angular* springs constraining the angles between pairs of the straight springs. Other configurations of springs are possible; the above choice serves our purposes well.




Figure 2: A two-dimensional spring system.

We then discretize the energy function, setting the total energy to

$$\sum_{b \in B} \int_b \psi(t) dt. \quad (4)$$

To use the above discretization of the energy function, we need the spring energy of each of these boxes to approximate $\int_b \psi(t) dt$, the mechanical energy for that box. This requires picking the appropriate spring constants for the two types of springs.

2.2.1 Straight Springs

To compute the constant k_1 for straight springs, we need a deformation whose energy comes only from the straight springs. Simple stretching, a type of deformation given by

$$\gamma(x, y, z) = ((a + 1)x, y, z) \quad (5)$$

is exactly that. It represents the stretching of the volume in the x -direction. We use the equations given in Section 2.1 to find the exact elastic energy E_b of a box b under simple stretching, obtaining

$$E_b = l_x l_y l_z \left(\frac{1}{2} \lambda + \mu \right) \left(\frac{1}{4} a^4 + a^3 + a^2 \right), \quad (6)$$

where l_x , l_y , and l_z are the undeformed spring lengths in the x , y and z directions.

To compute k_1 , we equate E_b with $\frac{1}{2} k_1 |l - l_x|^2$, which is the energy for a spring with stretched length l according to Hooke's law. We would now solve the equation for k_1 , except that doing so yields k_1 as a function of a . To obtain an actual constant we remove the a^3 and a^4 terms from E_b . Since the volumes we consider do not stretch beyond a small percentage of their undeformed dimension, a should be close to 0 and therefore a^3 and a^4 are small compared to a^2 . Solving the resulting equation yields k_1 and finally

$$E_{straight} = \frac{l_y l_z}{l_x} \left(\mu + \frac{\lambda}{2} \right) |l - l_x|^2. \quad (7)$$

The energies of the straight springs in the y and z directions are obtained similarly.




Figure 3: A deformed volume with minimal energy.

2.2.2 Angular Springs

An angular spring stretches with the cosine of the angle θ between the two straight springs it constrains. In other words, $\cos(\theta)$ is used as the distance in Hooke's law instead of a displacement.

Notice that angular springs have 0 energy when θ equals $\frac{\pi}{2}$, the undeformed position of the angle. Since $\sin(\theta) \approx \theta$ for small θ and $\sin(\frac{\pi}{2} - \theta) = \cos(\theta)$, $\cos(\theta) \approx \frac{\pi}{2} - \theta$ for θ close to $\frac{\pi}{2}$. To find a good spring constant k_2 for this type of spring, we use the shear deformation, which is:

$$\gamma(x, y, z) = (x, y + ax, z). \quad (8)$$

This deformation stretches our volume in addition to bending it. However, the deformation of the angular springs is much larger than that of the straight springs in this case so we ignore the latter. Once again, we find the exact energy of the deformation γ

$$E_b = \frac{1}{2} l_x l_y l_z \left(\frac{1}{4} \lambda a^4 + \frac{1}{2} \mu a^4 + \mu a^2 \right), \quad (9)$$

equate it with $\frac{1}{2} k_2 (\cos \theta)^2$, and solve for k_2 . Again we must drop the a^4 term to get a constant for k_2 . In this manner, we obtain

$$E_{angular} = \frac{1}{2} l_x l_y l_z \mu (\cos \theta)^2. \quad (10)$$

Because of energy minimization concerns (Section 2.4), we found it convenient to further restrict extreme bending. Since $-1 < \cos(\theta) < 1$, the energy contribution of the angular springs is bounded, whereas the energy of the straight springs is not. Because the energy provided by Eq. (10) is bounded, the volume might bend during energy minimization so that some angular spring approaches its maximum deformation. This is not good behavior, however, since a real volume should resist extreme bending, just as it resists extreme stretching. To correct this problem, we add an extra energy E_{extra} to $E_{angular}$ when $(\cos \theta)^2$ gets large, i.e. greater than some constant ζ . E_{extra} is given by

$$E_{extra} = \frac{(\cos \theta)^2 - \zeta}{1 - (\cos \theta)^2}. \quad (11)$$




Figure 4: At (a) the pipe is at an illegal configuration as its length is compressed. In (b) the minimization has restored the pipe close to its physical length. (c) shows how (a) was obtained given constraints at ends A and B .

2.3 Grasp/Manipulation Constraints

Now we impose manipulation constraints on our volume. That is, we restrict positions of some parts of V . In the spring-mass model, this means fixing the positions of some point masses. For example, we may want to hold a thick elastic plate on two sides and poke a third face (Figure 3). Then the positions of all the point masses on the surface of the two sides are fixed, and a point on the surface of the third face is fixed at a location translated from its undeformed position by the amount of poking.

2.4 Energy Minimization

In our planner we need to find deformations of minimal energy that satisfy given manipulation constraints. Because real objects tend toward minimal energy configurations, energy minimization approximates the shape that our volume physically assumes under the specified manipulation. We use a conjugate gradient descent method treating the energy as a scalar function of the positions of the point masses that are not fixed by the manipulation constraints. An illustration can be found in Figure 4 (a) and (b).

3 Path Planning

Main Algorithm We use a version of the Probabilistic Roadmap Planner (PRM) to plan paths for elastic volumes. For a detailed explanation of the scheme see [10]. Here we offer a brief outline of the algorithm.

We first generate a random manipulation constraint. Then, we compute the volume according to this constraint by minimizing the energy. We test if the obtained deformation abides by the elasticity limit (see below). If this deformation is valid, generate N random rigid-body transformations of this deformation, and then test each corresponding configuration for collision with the obstacles. The next step is to update the roadmap by trying to connect the new configurations generated to the configurations already

in the roadmap using the local planner. Finally, we enhance the roadmap by identifying “difficult” areas of the configuration space and generate M more nodes in those areas. We connect these nodes to the roadmap. After we have a dense enough roadmap, we search it to find a path from the initial to the goal configuration.

Local Planner An important issue in the above framework is the local planner. Here we follow [10] where the rigid-body transformation is computed first and then deformation paths are planned by interpolating between the manipulation constraints to form a sequence of intermediate constraints.

Elasticity Limit Our planner should ensure that all used deformations are within elasticity limits of the material so that our volume is not damaged during manipulation. The amount of deformation is defined by two coefficients:

- *Plane Strain Limit* This is how much the material stretches locally. To estimate this in our spring model, we find the linear spring which is most deformed.
- *Curvature Limit* We also do not want our volume to bend too much locally. We find the angular spring such that $\cos(\theta)/d$ is maximal, where d is the distance over which this spring approximates the curvature (e.g., $l_x + l_y$).

4 The Case of an Elastic Pipe

We now look at a specific type of volume, the case of an elastic pipe. Since our spring-mass representation uses a lattice, a long, thin rectangular solid represents the shape of the pipe (Figure 4(a) and (b)).

4.1 Manipulation Constraints

A manipulation constraint specifies the position and orientation of both ends of the pipe. We represent it with 5 parameters $(\bar{x}, \bar{y}, \bar{z}, \bar{\theta}, \bar{\varphi})$. If we imagine that the left end A

is held steady at a fixed orientation, then we can fix a coordinate frame whose origin coincides with A and whose x-axis follows along the length of the pipe at the left end. \bar{x} , \bar{y} , and \bar{z} refer to the position of the right end B with respect to the left one. The angles $\bar{\theta}$ and $\bar{\varphi}$ orient B with respect to the orientation of A . We do not allow twisting of the faces.

Initialization and Energy Minimization Given $(\bar{x}, \bar{y}, \bar{z}, \bar{\theta}, \bar{\varphi})$, we must then find a minimal energy deformation respecting these constraints. The conjugate-gradient method used in minimization requires an initial guess. Note that the initial position of the point masses in the lattice can effect the outcome and the speed of minimization. The initial guess needs to avoid unnatural starting conditions, which could cause the conjugate gradient method to become trapped in an undesirable high-energy local minimum. Furthermore, the method converges to a solution faster with an initial guess near a desirable minimum. A local minimum found starting from Figure 4(a) is shown in Figure 4(b).

The initial guess is specified as follows. We first construct the third degree polynomial curve which satisfies the following constraints:

- It is a parameterized curve P whose parameter ranges from 0 to 1.
- $P(0)$ and $P(1)$ are the center point masses of the left face A and the right face B , respectively.
- The tangent vector at $P(0)$ is the normal vector of the leftmost face, and the tangent vector at $P(1)$ is the normal vector of the rightmost face, scaled appropriately.

This gives a one dimensional approximation of the shape of the pipe (Figure 4(c)), although such an approximation becomes less accurate the closer A and B get to each other, since P does not try to preserve the pipe's length. To generate the full lattice we evenly divide the space of the parameter of P into as many divisions as we have point masses along the pipe's length. At each chosen point on the curve we place a cross section (square) of the pipe, perpendicular to the tangent of P at that point. The tangent vector uniquely specifies a plane, but the cross-sectional square could be rotated by any angle about the curve. To prevent twisting, we place the square in the plane by constructing it as a rotation of the left face in the following way.

- The tangent vector to the curve uniquely defines a normal plane. This plane is the image of the plane containing the left face by a rotation (or a translation if these planes are parallel). The axis of the rotation is the intersection between the planes.
- We apply the corresponding rotation (or translation) to the left face of the pipe to obtain a face in the normal plane. We translate this face in the normal plane to center it on the point of the curve.




Figure 5: Snapshots along a path of a deformable cable whose one end is fixed to a base.

This process provides us with a good initial guess for the positions of the lattice points of the pipe which has no twisting. However, if the manipulation constraints are outside a certain range, there may be self-intersections. We detect such occurrences and discard them.

4.2 Experimental Results

In Figure 5 the pipe remains fixed to a base, while it is manipulated around a frame and then attached to it. The pipe robot is represented by $32 \times 3 \times 3$ lattice of point masses. The parameters for the basic step in the PRM algorithm are $N = 20$ (random nodes), and $K = 5$. K is the number of neighbors considered for connection. Planning for this example took a mean time of 14.5 mins over 10 runs. All times reported in this paper were gathered on an SGI R10000.

The example in Figure 6 involves planning an animation to take an elastic pipe through an L-shaped hole inside a cube. Here the pipe is $21 \times 3 \times 3$, with $N = 200$ and $K = 20$. The path took 8h 39min, considerably longer than the above, since only a small portion of the possible deformations will fit through the obstacle.

A crucial part of our algorithm is the energy minimiza-

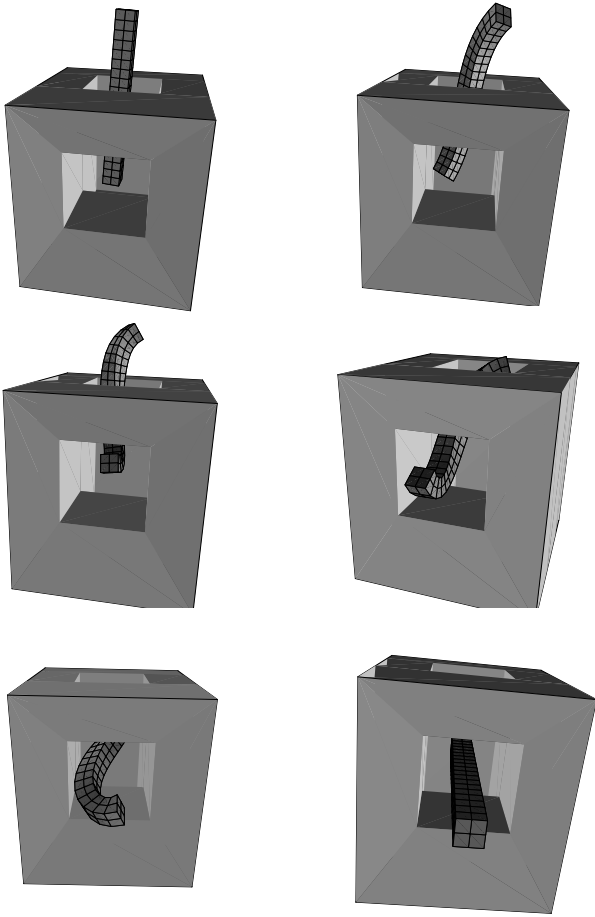


Figure 6: Snapshots of an animation of a deformable pipe.

tion. With the pipe from the above example, we selected and minimized 100 random deformations. The mean time was 1.12s, with a standard deviation of .98.

5 Discussion

The paper developed a simple model for representing three dimensional deformable objects in the context of path planning. We demonstrated the validity of our approach by applying it to elastic pipes. Our work raises several interesting issues. These include representing more complex shapes with lattices of springs, finding more efficient energy minimization procedures, testing the limits of the current geometric model ([18] is especially interesting), allowing more complex manipulation constraints and contact with the obstacles in the environment, and smoothing the resulting paths to avoid unnecessary deformations.

Acknowledgments Work on this paper by Scott Owens and Elliot Anshelevich has been supported by IRI-970228 and the corresponding REU portion of this award. Work on this paper by Lydia Kavraki and Florent Lamiraux has been supported in part by NSF CAREER Award IRI-970228 and NSF CISE SA1728-21122N. The authors would like to thank the members of the Physical Computing Group at Rice for their comments.

References

- [1] D. Baraff and A. Witkin. Large steps in cloth simulation. *Computer Graphics (SIGGRAPH'98)*, pages 303–308, 1998.
- [2] H. Chang and T. Li. Assembly maintainability study with motion planning. In *Proc. IEEE Int. Conf. on Rob. and Autom.*, pages 1012–1019, 1995.
- [3] S. Cotin, H. Delingette, and N. Ayache. Real-time elastic deformations of soft tissues for surgery simulation. *IEEE Transactions on Visualization and Computer Graphics*, pages 62–73, 1999.
- [4] S. Gibson and B. Mirtich. A survey of deformable modeling in computer graphics. Technical Report TR-97-19, MERL, 1997.
- [5] K. Gupta and A. P. del Pobil. *Practical Motion Planning in Robotics*. John Wiley, West Sussex, England, 1998.
- [6] D. James and D. Pai. Accurate real time deformable objects. In *Proc. of SIGGRAPH'99*, 1999.
- [7] A. Joukhadar and C. Laugier. Dynamic simulation: Model, basic algorithms, and optimization. In J.-P. Laumond and M. Overmars, editors, *Algorithms for Robotic Motion and Manipulation*, pages 419–434. A. K. Peters, 1997.
- [8] L. Kavraki and J.-C. Latombe. Randomized preprocessing of configuration space for fast path planning. In *Proc. IEEE Int. Conf. Robotics and Automation*, pages 2138–2145, San Diego, CA, 1994.
- [9] L. E. Kavraki, F. Lamiraux, and C. Holleman. Towards planning for elastic objects. In P. Agarwal, L. E. Kavraki, and M. Mason, editors, *Robotics: The Algorithmic Perspective*, pages 313–326. AK Peters, Natick, MA, 1998.
- [10] F. Lamiraux and L. E. Kavraki. Path planning for elastic plates under manipulation constraints. In *Int. Conf. on Robotics and Automation*, Detroit, MI, 1999.
- [11] H. Nakagaki and K. Kitagaki. Study of deformation tasks of a flexible wire. In *Proc. IEEE Int. Conf. on Rob. and Autom.*, Albuquerque, NM, 1997.
- [12] E. Promayon, P. Baconnier, and C. Puech. Physically-based deformations constrained in displacements and volume. *EUROGRAPHICS*, 1996.
- [13] X. Provot. Deformation constraints in a mass-spring model to describe rigid cloth behavior. In *Proc. Graphics Interface '95*, pages 147–154, 1995.
- [14] A. Singh, J.-C. Latombe, and D. Brutlag. Path planning for molecular docking. Personal Communication.
- [15] D. Terzopoulos, J. Platt, A. Barr, and K. Fleischer. Elastically deformable models. *Computer Graphics (SIGGRAPH'87)*, 21(4):205–214, 1987.
- [16] T. van Walsum and M. Viergever. Deformable b-splines for catheter simulation. Technical report, Image Science Institute, Utrecht University, 1998.
- [17] G. Wempner. *Mechanics of Solids with applications to thin bodies*. McGraw-Hill, NY, 1991.
- [18] Y. Zhuang and J. Canny. Real-time and physically realistic simulation of global deformation. In *SIGGRAPH99 Sketches and Applications*, 1999.