

Toward Automated Tissue Retraction in Robot-Assisted Surgery

Sachin Patil and Ron Alterovitz

Abstract—Robotic surgical assistants are enhancing physician performance, enabling physicians to perform more delicate and precise minimally invasive surgery. However, these devices are currently tele-operated and lack autonomy. In this paper, we present initial steps toward automating a commonly performed surgical task, *tissue retraction*, which involves grasping and lifting a thin layer of tissue to expose an underlying area. Given a model of tissues in the vicinity, our method computes a motion plan for a 6-DOF gripper that grasps a tissue flap at an optimal location and retracts it such that an underlying target is fully visible. The planner considers three optimization objectives relevant to medical applications: minimizing the maximum deformation energy, minimizing maximum stress, and minimizing the control effort in lifting the tissue flap. The planner can be used to locally improve physician specified retraction trajectories based on the optimization criteria or to compute a de novo plan. We use a physically-based simulation to compute equilibrium configurations of the tissue flap subject to manipulation constraints. These configurations are used with a sampling-based planner to explore the space of deformations and compute an optimal plan subject to discretization and modeling error. Our experimental results illustrate the ability of the method to compute retractions for heterogeneous tissues while avoiding obstacles and minimizing tissue damage.

I. INTRODUCTION

Robotic surgical assistants (RSAs), such as Intuitive Surgical's da Vinci[®] system [8], are enhancing physician performance, enabling them to perform more delicate and precise minimally invasive surgery. Studies have shown that these robots can improve procedure success rates, reduce bleeding, and decrease recovery time [9]. As a result, hospitals are increasingly adopting RSAs, with over 1,000 installed worldwide [8]. In their current form, RSAs are tele-operated; the surgeon performs the medical procedure using an input device outside the patient, and the robotic device duplicates the motions of the input device inside the patient (possibly at a different scale and with smoothed motions). Because of the dependence on tele-operation, clinically-used RSAs currently lack autonomy and require direct control by physicians.

Integrating motion planning with RSAs has the potential to enable RSAs to perform certain motions autonomously or to improve physician-specified trajectories to minimize tissue stress and forces. In this paper, we focus on a relatively simple but commonly performed task: tissue retraction. The goal is to manipulate an outer layer of tissue to provide the physician with a line of sight to an area of interest underneath while avoiding contact with obstacles and nearby sensitive structures (see Fig. 1). The manipulation must balance competing objectives: provide sufficient exposure and avoid excessively large forces that damage the tissues.

S. Patil and R. Alterovitz are with the Department of Computer Science, University of North Carolina at Chapel Hill, USA
sachin@cs.unc.edu, ron@cs.unc.edu




Fig. 1. During tissue retraction, the tissue flap (initially flat) is manipulated by the gripper such that the area of interest (marked by the red circle) is completely visible from a given camera viewpoint (left frame). An obstacle (protruding from left in left frame, shown at bottom from a different perspective in right frame) must be avoided during retraction.

Enabling RSAs to autonomously perform sub-procedures like tissue retraction could have a significant impact on patient care. Adding some level of autonomy has the potential to reduce surgical errors by enabling the physician to focus on the important, challenging aspects of a procedure rather than being distracted by motion control of devices. This would also enable RSAs to utilize more than two manipulators simultaneously, enabling greater dexterity and faster procedure completion times. Also, the absence of force-feedback in most current tele-operated surgical systems hampers physicians' ability to correctly estimate crucial physical quantities such as exerted forces and stress buildups in the deformed tissue. Underestimating these quantities during a procedure may cause tissue damage, resulting in side effects and longer recovery times. RSAs utilizing a motion planner with an embedded physically-based model could improve physician-specified trajectories to minimize unnecessary tissue damage.

We present an approach to compute plans for automated tissue retraction using a generic 6-DOF gripper. Our method can either be used for fully automating the tissue retraction task (global optimization) or for improving physician specified retraction trajectories (local optimization). We consider three optimization objectives relevant to medical applications: minimizing the maximum deformation energy, minimizing maximum tissue stress, and minimizing the cumulative control effort. The method relies on a physically-based simulation to compute equilibrium states of the tissue flap for sampled gripper configurations. These states are used to construct a global roadmap to explore the space of deformations and compute an optimal plan that avoids obstacles. To compute the portion of a target area visible beneath the tissue flap, we use a fast visibility test using occlusion queries on a graphics processing unit (GPU). Our planner computes trajectories for the gripper to accomplish the tissue retraction task while avoiding obstacles, as shown in Fig. 1.

This planner is an initial step toward automating the task of tissue retraction during robot-assisted surgery. Ultimately, we plan to improve planner performance and integrate it with pre-operative medical imaging and RSA hardware.

II. RELATED WORK

Robotic Surgical Assistants (RSAs): Tele-robotic or master/slave systems, such as the da Vinci[®] system [8], operate under direct control of the surgeon and offer improved precision and dexterity for minimally invasive procedures compared to non-robotic approaches [8], [15]. Some autonomous systems have been successfully used for performing surgeries on non-deformable tissue (such as bones) when detailed quantitative pre-operative plans of the procedure are available [29]. Our approach proposes to autonomously perform the task of tissue retraction involving deformable tissues.

Modeling of Deformable Objects: Physically-based simulation of deformable objects is a well-studied area in solid mechanics [30] and computer graphics [22]. The choice of the underlying model is application-specific and has a substantial influence on the physical accuracy of the estimated deformations. Finite Element Methods (FEM) are based on the theory of continuum mechanics [30] and simulate deformations accurately at the cost of increased computational complexity when model parameters are known. Nonlinear FEM models are preferred when the correctness of the simulation is important and have been successfully used for simulating soft tissue for surgical simulations [6], [24], [23].

Motion Planning/Manipulation of Deformable Objects: Research in robot motion planning and manipulation has historically focused on robots composed of rigid links operating in environments with rigid objects [20]. Recent work has begun to explore motion planning for deformable robots in static environments. Lamiroux et al. [19] developed the f-PRM framework to plan paths for flexible robots of simple geometric shapes such as surface patches or simple volumetric elements. Bayazit et al. [3] use a two-tier approach based on probabilistic roadmaps (PRM) and free-form deformation (FFD) to plan paths for deformable robots. Gayle et al. [12] use a constraint-based motion planner for deformable objects modeled as mass-spring systems. Rodriguez et al. [25] use rapidly exploring random trees (RRT) to plan for robots in completely deformable environments. Frank et al. [11] use co-rotational linear FEM with a PRM-based planner to achieve significant speedups for path planning in deformable environments. Motion planning algorithms have also been developed for clinical applications, including deformable robots traveling through body cavities [12], deformable linear objects [21], [26], and flexible needle devices traveling through deformable tissue [2].

Prior work on robot grasping and manipulation have generally assumed that the grasped object is rigid [4]. Howard et al. [17] model deformable parts using a mass-spring model and use a neural network to control the gripper. Hirai et al. [16] propose a robust control law for local manipulations of deformable parts using tactile and video feedback. Gopalakrishnan and Goldberg [13] extend the

form closure framework to grasping of planar, deformable objects using a two-point gripper. We address the problem of motion planning for constrained deformable objects subject to contact and manipulation constraints.

In concurrent work to ours, Jansen et al. [18] combine the D-Space framework proposed by Gopalakrishnan and Goldberg [13] with a geometric approach to compute candidate grasp locations and optimal trajectories for automated tissue retraction. This computationally fast approach considers a 2D cross-sectional model of the tissue flap and computes trajectories for the gripper in a plane. In contrast, our approach works with a fully 3D tissue model, handles generic contact and manipulation constraints, and searches for optimal solutions in constrained environments containing obstacles that must be avoided.

III. PROBLEM STATEMENT

We assume that the geometry of the tissue, including the flap to be retracted and the obstacles, is known from data obtained from medical images or other sensors. We model the tissue as a 3D deformable body \mathcal{M} (represented as a tetrahedral finite element mesh \mathcal{M}_{ref} for simulation purposes). We assume that \mathcal{M} has known material properties, which may be heterogenous throughout the volume of the tissue. We model the robot gripper as a rigid 6-DOF gripper with a non-zero surface area of contact. Let \mathbf{q}_i be a configuration of the gripper in $SE(3)$.

We do not plan for the motion of the actuators or the linkages but only consider their cumulative effect on the gripper pose. We assume that the tissue moves slowly as it is manipulated by the gripper such that dynamics can be ignored and we can approximate the process as quasi-static. We also assume that the motion of the gripper and the subsequent tissue deformations are deterministic; we plan to consider the effects of uncertainty due to factors such as material parameters, actuation, and slippage in future work.

Exerting excessive forces at the gripper or inducing excessively large deformations of soft tissue during manipulation can cause tissue damage (such as tearing of tissue). We propose the following optimization objectives which quantify these crucial physical quantities:

- **Minimizing maximum deformation energy:** The total deformation energy of a deformed body provides a quantitative measure of the extent of its overall deformation. In a clinical application, paths that involve high-energy intermediate states may introduce plastic deformations and cause other irreversible tissue damage. Paths with lower deformation energies will be less susceptible to these effects.
- **Minimizing maximum stress:** The maximum stress (tensile or compressive) encountered in the deformed body represents an upper limit to the forces that can be applied to deformable objects without causing damage (such as tearing and fracture) [30]. Minimizing the maximum stress helps discard states that are close or beyond the point of damage, resulting in safe plans.

- **Minimizing total control effort:** We define the total control effort expended in performing the task as a combination of the total deformation energy and the gripper displacement, which measures mechanical effort. Minimizing the total control effort balances two competing objectives: minimizing the deformation energy of the intermediate states along the path and minimizing the total mechanical effort involved in executing the plan.

The choice of optimization objective depends on several factors such as maneuverability of the gripper in the workspace and the type of tissue involved in the retraction task and should be determined by the physician during the procedure.

The problem can now be broken down into two subproblems and formally stated as follows:

Problem 1: To compute a sequence of gripper configurations (position and orientation) that optimizes the chosen objective subject to the constraint that the area of interest beneath the tissue flap becomes completely visible from the physician’s fixed viewpoint.

Input: Initial configuration of the gripper \mathbf{q}_0 , tissue mesh in its reference state \mathcal{M}_{ref} , area of interest beneath the tissue flap \mathcal{A} , fixed camera viewpoint \mathcal{V} , description of the environment (set of obstacles \mathcal{O}) and tissue material parameters (e.g. Young’s modulus E , Poisson’s ratio ν , and density ρ).

Output: An optimal trajectory (discrete sequence of configurations) $\mathcal{P} = (\mathbf{q}_0, \mathbf{q}_1, \dots, \mathbf{q}_n)$ of the gripper to manipulate the tissue flap and perform the retraction procedure.

Problem 2: The first problem assumes that the initial grasp location of the gripper is known. This can be generalized to also compute the optimal grasp location for the gripper \mathbf{q}_0^* , which optimizes the chosen objective over all possible grasp locations for a given scenario.

IV. MANIPULATION AND PATH PLANNING

The two key components of our approach are a sampling-based global planner and a physically-based simulator as shown in Fig. 2. For each randomly sampled configuration of the gripper $\mathbf{q}_i \in SE(3)$ generated by the planner, the simulator computes the corresponding deformed equilibrium state of the mesh \mathcal{M}_i subject to contact and manipulation constraints (see Sec. V). These configurations are used to explore the space of deformations and compute an optimal plan. It should be noted that even though the planner is essentially independent of the simulator, the accuracy, correctness, and speed of computation of the results is implicitly dependent on the chosen simulator.

A. Generating Valid Samples

Manipulation Constraints: The tissue flap in the mesh \mathcal{M} is assumed to be anchored at a subset of nodes (usually along one or more edges of the flap) and can be manipulated by specifying a non-zero displacement of the grasped nodes along the unconstrained edges of the flap boundary. To account for a realistic grip, we constrain multiple nodes on the surface of the flap at the gripper location. Based on the initial grasp location in the reference state \mathcal{M}_{ref} , we define all nodes within the bounding box of the gripper jaws as

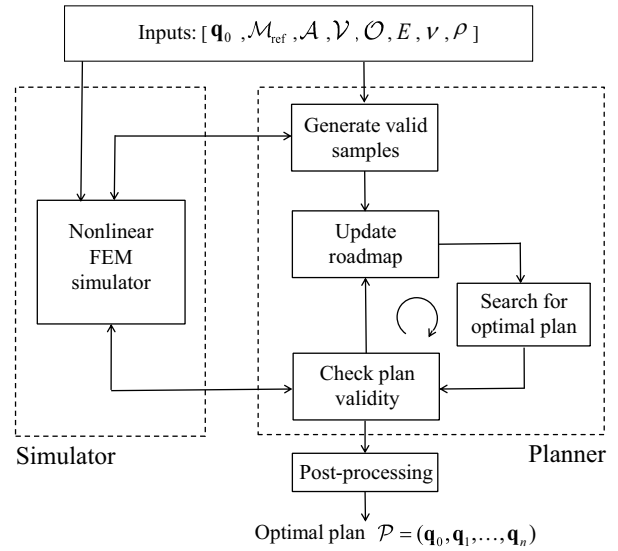


Fig. 2. System Overview: The probabilistic roadmap (PRM) based global planner (Planner) relies on the deformable body simulator (Simulator) to compute static, equilibrium states of the tissue flap corresponding to randomly sampled gripper configurations and generate an initial roadmap. The roadmap is queried for an optimal plan based on the chosen optimization objective, which is validated using the Simulator using a lazy approach. Notation for input/output explained in Sec. III.

“grasped” nodes. We generate manipulation constraints of the tissue by randomly sampling configurations $\mathbf{q}_i \in SE(3)$ of the 6-DOF gripper. For a given sample \mathbf{q}_i , we displace and fix all grasped nodes by the transformation given by \mathbf{q}_i . As described in Sec. V, we can then compute an equilibrium mesh state \mathcal{M}_i corresponding to the gripper configuration \mathbf{q}_i . **Checking Mesh Validity:** We consider a mesh state \mathcal{M}_i (and corresponding configuration \mathbf{q}_i) to be *valid* if the final topology of the mesh is valid (i.e. no inverted elements or self-penetration). We use Tetgen [28] to discretize the mesh and reliably detect self-penetration.

Sample Optimization: Since the gripper configuration \mathbf{q}_i is sampled randomly, the equilibrium mesh states may have undesirably high stresses in the elements constrained by the gripper. For every valid sample generated, we perform local gradient-descent optimization on the orientation of the gripper (while keeping its position in \mathbb{R}^3 constant) to alleviate the stress on the constrained elements (as shown in Fig. 3). Our experiments indicate that fewer states are eventually required in the roadmap to compute a feasible plan if this local improvement is applied.

Local Improvement of Physician-Specified Trajectories: In the absence of force-feedback mechanisms and/or restricted visibility, our method can also be used for local improvement of physician specified retraction trajectories in challenging scenarios. The physician indicates an approximate retraction path for the gripper (using the tele-operated surgical system). In our implementation, the physician indicates a final gripper position (in \mathbb{R}^3) and the trajectory is implicitly defined as a straight line joining the initial and user-specified gripper position (as shown in Fig. 4). A multi-linear trajectory (based on multiple way-points) can also be considered. The gripper configurations \mathbf{q}_i are then uniformly randomly sampled in




Fig. 3. For each sampled gripper configuration, we optimize orientation to minimize the stresses in the elements constrained by the gripper. For a randomly generated configuration, these stresses can be high (left). Local optimization re-orientes the gripper to minimize these stresses (right).

a user-defined neighborhood of this trajectory. It should be noted that the quality of the final solution obtained will depend on the physician-specified trajectory and the choice of neighborhood sample distribution.

B. Computing Visibility

The objective of the retraction task is to expose an area of interest \mathcal{A} from the given camera viewpoint \mathcal{V} . A mesh state \mathcal{M}_i is designed as a goal state if \mathcal{A} is completely visible from the given viewpoint.

There are two primary types of methods to compute visibility. Object-space methods intersect the view frustum of the camera with primitives in the scene and compute the exact area visible from the given viewpoint. These methods are exact but computationally expensive. Image-space methods render a scene as an image from the camera viewpoint and compute the area visible in terms of pixels. These methods are approximate but fast. The pixel accuracy suffices since we are only interested in the exposure of a pre-defined area of interest. For highly-parallel, fast performance, we implemented the visibility computation using occlusion queries on graphics processing units (GPUs) [7]. Based on results of the visibility tests, we identify the set of goal mesh states $\mathbf{M}_g : \{\mathcal{M}_g^{(1)}, \dots, \mathcal{M}_g^{(m)}\}$ and corresponding gripper configurations $\mathbf{Q}_g : \{\mathbf{q}_g^{(1)}, \dots, \mathbf{q}_g^{(m)}\}$.

C. Constructing the Roadmap

A PRM-based planner can be used to compute a plan from the start configuration \mathbf{q}_s and a goal configuration from set \mathbf{Q}_g . We store the mesh state \mathcal{M}_i along with each sample \mathbf{q}_i for computing the distance metric defined below.

Distance Measure: We use a heuristic distance measure [19] for selecting the nearest neighbors of a sample. If \mathcal{M}_i and \mathcal{M}_j are the mesh equilibrium states corresponding to configurations \mathbf{q}_i and \mathbf{q}_j , then the distance measure is defined as $d(\mathbf{q}_i, \mathbf{q}_j) = d_{\mathcal{D}}(\mathcal{M}_i, \mathcal{M}_j) + d_{\mathcal{C}}(\mathbf{q}_i, \mathbf{q}_j)$, where $d_{\mathcal{D}}(\mathcal{M}_i, \mathcal{M}_j)$ is the DISP distance metric defined as the maximum length over all displacement vectors for each node on the surface of \mathcal{M}_i to \mathcal{M}_j and $d_{\mathcal{C}}(\mathbf{q}_i, \mathbf{q}_j)$ is the distance metric between two configurations defined for the $SE(3)$ group [20]. Attempts to weight $d_{\mathcal{D}}$ and $d_{\mathcal{C}}$ have not yielded better results but we found that using only $d_{\mathcal{D}}$ works reasonably well.

Local Planning: We define a path between two samples represented by gripper configurations \mathbf{q}_i and \mathbf{q}_j by linearly




Fig. 4. Heuristic sampling strategy for local improvement of physician specified retraction trajectories. In our test cases, the user specifies a trajectory by an initial gripper location and a final gripper position in \mathbb{R}^3 (left). The gripper configurations \mathbf{q}_i are sampled within a user-defined neighborhood of the trajectory (right).

interpolating between them in $SE(3)$ [20] and executing the simulator to obtain intermediate mesh states. The edge connecting any two samples in the roadmap is considered to be valid if each of the intermediate mesh states is valid (see Sec. IV-A) and if both the mesh and the gripper do not collide with any obstacles in the workspace. We use PQP [14] for collision detection with obstacles.

D. Assigning Roadmap Edge Weights

We use the optimization objectives outlined in Sec. III to assign weights to edges in the roadmap. These weights are computed as follows:

Minimizing maximum deformation energy: The cost of an edge is set to be the maximum deformation energy encountered as the gripper moves from configuration \mathbf{q}_i to \mathbf{q}_j . The total energy of a single tetrahedral element τ is given by the sum of the internal elastic energy stored in the tetrahedron and the work done by the applied external forces \mathbf{f}_{ext} on each of the nodes $\eta \in \tau$. The total deformation energy Π_i of the mesh state \mathcal{M}_i is: $\Pi_i = \sum_{\tau \in \mathcal{M}_i} v(\tau) \cdot W(\mathbf{E}) + \sum_{\eta} \mathbf{f}_{\text{ext}}^{\eta} \cdot \mathbf{u}_{\eta}$, where $v(\tau)$ is volume of the tetrahedron τ , \mathbf{u}_{η} is displacement of node η , \mathbf{E} is the Green strain tensor, and $W(\mathbf{E})$ is the energy density function for the chosen material model. We refer the reader to [30], [23] for additional details.

Minimizing maximum stress: The cost of an edge is chosen to be the maximum stress (tensile or compressive) encountered as the gripper moves from \mathbf{q}_i and \mathbf{q}_j . For a tetrahedron $\tau \in \mathcal{M}_i$, let $\sigma^i(\mathbf{S}_{\tau}), i \in \{1, 2, 3\}$ be the i th eigenvalue of the second Piola stress tensor \mathbf{S}_{τ} . Positive eigenvalues correspond to tensile stresses and negative ones to compressive stresses. Since \mathbf{S}_{τ} is real and symmetric, it will have three real, not necessarily unique, eigenvalues. The maximum stress in τ is then given by $\max(\text{abs}(\sigma^i)), i \in \{1, 2, 3\}$ [30].

Minimizing control effort in lifting tissue flap: We express the total control effort C using the metric defined in Moll et al. [21] along a specified path $(\mathbf{q}_1, \dots, \mathbf{q}_n)$ as: $C = \sum_{i=1}^{n-1} \frac{1}{2} d(\mathbf{q}_i, \mathbf{q}_{i+1}) (\Pi_i + \Pi_{i+1})$. A better approximation of this integral can be obtained by considering a larger number of subdivisions along the path.

E. Computing Optimal Plans

The local planner requires running the simulation along the interpolated path, which is computationally expensive. We use a *lazy* local planning scheme [5] to reduce computation

times. We construct an initial roadmap by assigning edge weights using configurations at endpoints only. For optimization objectives such as minimizing maximum deformation energy or tissue stress, we set the edge weight to be the maximum of the objective function value at the two endpoints (which serves as a lower bound for the actual edge weight). For minimizing the total control effort, the actual edge weight is readily computed.

We search for a minimum-cost path from \mathbf{q}_s to \mathbf{Q}_g in this roadmap. We use Dijkstra’s shortest path algorithm for searching for an optimal path from \mathbf{q}_0 to \mathbf{Q}_g for additive optimization objectives (such as minimizing the total control effort). The problem of minimizing the maximum edge weight in an acyclic graph (the edge weight is the deformation energy or tissue stress in our case) is referred to as the Bottleneck Shortest Paths problem and can be solved by a minor modification to Dijkstra’s algorithm [1].

The local planner is used to check if all the edges along the path are valid. Any invalid edges are discarded and valid edges are marked as processed and their edge weights updated. This process is repeated until a valid (optimal) plan is obtained or no solution is found. As the number of samples generated in the pre-processing phase increase, the solution converges to the true optimal.

Since the paths generated by PRM-like methods are not necessarily short or smooth, we use an iterative shortening and smoothing scheme to improve the final solution. Given a solution path comprising of discrete gripper configurations, $\mathcal{P} = (\mathbf{q}_1, \dots, \mathbf{q}_n)$, the final continuous sequence of controls is obtained by interpolating between the configurations using a cubic spline [10] to ensure C^1 -continuity.

V. SIMULATION OF TISSUE FLAP RETRACTION

Simulating soft tissue is challenging since simulation must provide realistic response to manipulations such as grasping and pulling. We chose to model tissue deformations using a finite element method (FEM) and nonlinear continuum mechanics to account for large tissue deformations. The tissue itself is assumed to be isotropic and nearly incompressible and exhibit hyper-elasticity according to the St. Venant Kirchhoff material model (commonly chosen to represent bio-mechanical deformable objects [24], [23]). We refer the reader to [30] for an introduction to nonlinear FEM.

The equilibrium configuration of the tissue mesh is determined by balancing all global external forces \mathbf{f}_{ext} with all internal elastic forces $\Phi(\mathbf{u})$, where \mathbf{u} represents the global displacement of all the nodes \mathcal{N} in the mesh and Φ is a (possibly nonlinear) function describing the internal elastic forces as a function of the nodal displacements \mathbf{u} .

Assembling External Forces: The tissue flap is subject to external forces exerted by gravity as well as other contact and manipulation constraints. To compute gravitational forces, we use a lumped mass formulation [24]. We resolve collisions of the flap with the underlying tissue (modeled as a plane in our experiments) using a penalty-based scheme where reaction forces are proportional to the penetration depth of the intersecting nodes. More sophisticated collision resolution

schemes could also be used [11]. We do not directly consider the interaction forces between the gripper and the tissue flap and instead fix tissue nodes inside the gripper to the gripper using displacement constraints.

Assembling Elastic Forces: The elastic force acting on a node is given by the negative gradient of the elastic energy in the element with respect to the nodal displacement. The total elastic force acting on a node is obtained by summing up the elastic forces exerted by all the individual tetrahedra that are incident to the node. We follow the approach suggested in [23] for computing elastic forces. We also add penalty forces to each node of the tetrahedron proportional to the variation in the volume of the tetrahedron as suggested in [24]. This allows us to model the nearly incompressible nature of soft tissue and to prevent inversion of tetrahedra under manipulation constraints.

Solving for Static Equilibrium: The tissue flap is anchored at a subset of nodes (usually along one or more edges along the boundary of the flap) and is also manipulated by the gripper. Let \mathcal{N}_{fix} denote the set of *anchored* nodes (which have zero displacement) and *grasped* nodes (for which the displacement is determined by the gripper configuration \mathbf{q}_i). The equilibrium state of the mesh \mathcal{M}_i corresponding to the gripper configuration \mathbf{q}_i is given by solving the (reduced) system of equations, $\tilde{\Phi}(\tilde{\mathbf{u}}) = \tilde{\mathbf{f}}_{\text{ext}}$, where $\tilde{\mathbf{u}}$ represents the global displacement of all the nodes in the set $\mathcal{N} \setminus \mathcal{N}_{\text{fix}}$ and $\tilde{\mathbf{f}}_{\text{ext}}$ represents the external forces acting on the corresponding nodes. We solve this system of equations using an iterative nonlinear conjugate gradient solver [27].

VI. RESULTS

We implemented the planner in C++ using OpenMP for parallelization and tested it on a 3.33 GHz 4-core Intel® i7™ PC. We use a rectangular tissue flap model of dimensions 5 cm \times 5 cm \times 0.25 cm for our experiments. The tetrahedralized tissue flap mesh (constructed using Tetgen [28]) contains 4000 elements. We set tissue density $\rho = 1000$ kg/m³, Poisson’s ratio $\nu = 0.45$, and Young’s Modulus $E = 1$ MPa (unless otherwise specified) and gravity at 9.8 m/s² acting downwards.

We applied our method in two contexts. The first is full autonomy, where the planner globally sampled the entire feasible space to compute a de novo plan. The second is local improvement of a physician-specified trajectory, as shown in Fig. 4. We evaluated our method using two tissue retraction scenarios described below.

Scenario 1: Fig. 5 shows an incision in the tissue from an overhead viewpoint. The objective here is to expose the arbitrarily shaped area of interest by parting the two tissue flaps while minimizing the maximum deformation energy and avoiding obstacles. Since the two tissue flaps are symmetric in the initial boundary conditions (each tissue flap has one straight non-fixed edge), we compute equilibrium mesh states in the pre-processing stage for a single tissue flap. The optimal initial grasp configuration \mathbf{q}_0^* is obtained by performing an exhaustive search over a set of 5 possible initial grasp configurations along the free edge of each tissue

flap. The trajectories of the grippers are then computed by our framework independently for the two tissue flaps. Fig. 5 shows the optimal initial grasp positions and the retraction trajectories of the grippers.

Scenario 2: Fig. 6 shows a scenario in which the tissue flap is free along two straight edges and the gripper is allowed to grasp anywhere along these edges. An obstacle protrudes from the side and extends over the tissue flap. The tissue flap is modeled as a heterogeneous structure with vessels that are stiffer in comparison to the surrounding tissue. In our implementation, we use a threshold image of the tissue texture to segment the vessels and assign appropriate tissue parameters to the mesh elements in the flap. In this experiment, we set the Young’s Modulus for the vessels to be 1 MPa and for the surrounding tissue to 250 KPa. The optimal initial grasp configuration \mathbf{q}_0^* is obtained by performing an exhaustive search over a set of 9 possible initial grasp configurations. Figs. 6 and 7 show the optimal grasp locations and retraction trajectories for the gripper when the objective is to minimize the control effort and the tissue stress respectively. In both cases, the gripper successfully lifts the tissue while avoiding obstacle intersection. Fig. 4 shows the optimized retraction trajectory obtained after local improvement of a physician-specified tissue retraction trajectory.

Performance: Table I provides detailed computation times. Table II compares the variation in the chosen optimization objective for the global optimization method to compute a de novo plan versus a local improvement method.

Scenario	Num samples	Sample generation (hrs)	Lazy planning			Total (secs)
			Collision detection (secs)	Mesh validation (secs)	FEM sim (secs)	
1	1000	1.21 (\pm 0.15)	0.56	24.94	7.11	32.71 (\pm 5.98)
2	1000	3.03 (\pm 0.23)	1.80	32.85	48.15	83.32 (\pm 38.98)

TABLE I

METHOD PERFORMANCE. AVERAGE TIMES AND STANDARD DEVIATIONS COMPUTED OVER MULTIPLE INITIAL GRASP LOCATIONS $\mathbf{q}_0^{(i)}$.

Scenario	Global optimization		Local improvement		% RMS error	
	Num samples	Sample generation (hrs)	Num samples	Sample generation (hrs)	Min. maximum energy	Min. maximum stress
1	1000	1.21 (\pm 0.15)	200	0.16 (\pm 0.01)	88.9%	73.6%
2	1000	3.03 (\pm 0.23)	200	0.29 (\pm 0.02)	65.6%	90.9%

TABLE II

COMPARISON OF THE GLOBAL OPTIMIZATION AND LOCAL IMPROVEMENT METHODS. % RMS ERROR IN COMPUTED FOR LOCAL IMPROVEMENT AS COMPARED TO GLOBAL OPTIMIZATION OVER 10 SUCCESSFUL RUNS FOR THE OPTIMAL GRIPPER CONFIGURATION \mathbf{q}_0^* .

The sample generation process involves expensive simulations and is the computational bottleneck of our method. We plan to build on recent advances in finite element simulation [6] to enable interactive execution of our method. Nevertheless, our method is successfully able to compute retractions for heterogeneous tissues in constrained environments, while avoiding obstacles and minimizing tissue damage.

VII. CONCLUSION

We have described a framework that takes steps toward automating the task of tissue retraction in robot-assisted surgical

procedures. The method is directly applicable to arbitrarily shaped tissue flaps. Given a model of the tissue flap, a physically-based simulator generates equilibrium mesh states corresponding to randomly sampled gripper configurations. These can be used with a sampling-based motion planner to compute an optimal sequence of controls for the gripper. We demonstrated the method for several optimization objectives particularly relevant to medical applications: minimizing the maximum deformation energy, minimizing maximum stress, and minimizing the control effort in lifting the tissue flap.

A key challenge in terms of future work is to improve the speed of the method to enable interactive execution in a clinical setting. Another important issue that needs to be addressed is the discrepancy between simulation and observed tissue behavior due to uncertainty in material parameters, gripper actuation, and tool-tissue interaction. As in prior work on Stochastic Motion Roadmaps [2], we will investigate approaches to encode uncertainty information into the global roadmap used for tissue retraction planning.

VIII. ACKNOWLEDGEMENTS

This research was supported in part by the National Science Foundation (NSF) grant #IIS-0905344. The authors thank M. Cenk Cavusoglu, Ken Goldberg, Wyatt Newman, and Pieter Abbeel for their valuable input on this problem.

REFERENCES

- [1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms and Applications*. Prentice Hall, 1993.
- [2] R. Alterovitz, T. Simeon, and K. Goldberg, “The stochastic motion roadmap: A sampling framework for planning with Markov motion uncertainty,” in *Proc. Robotics: Science and Systems (RSS)*, 2007.
- [3] O. B. Bayazit, J.-M. Lien, and N. M. Amato, “Probabilistic roadmap motion planning for deformable objects,” in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2002, pp. 2126–2133.
- [4] A. Bicchi and V. Kumar, “Robotic grasping and contact: A review,” in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2000, pp. 348–353.
- [5] R. Bohlin and L. E. Kavraki, “Path planning using lazy prm,” in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2000, pp. 521–528.
- [6] N. Chentanez, *et al.*, “Interactive simulation of surgical needle insertion and steering,” *ACM Transactions on Graphics (Proc. SIGGRAPH)*, vol. 28, no. 3, pp. 88:1–88:10, Aug. 2009.
- [7] M. Craighead and D. Ginsburg, “ARB occlusion query,” 2001. [Online]. Available: http://oss.sgi.com/projects/ogl-sample/registry/ARB/occlusion_query.txt
- [8] da Vinci[®] Surgical System, “<http://www.intuitivesurgical.com/>,” 2005.
- [9] S. A. Darzi and Y. Munz, “The impact of minimally invasive surgical techniques,” in *Annu Rev Med.*, vol. 55, 2004, pp. 223–237.
- [10] M. DeLoura, *Game Programming Gems 2*. Charles River Media, 2002.
- [11] B. Frank, M. Becker, C. Stachniss, W. Burgard, and M. Teschner, “Efficient path planning for mobile robots in environments with deformable objects,” in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2008, pp. 3737–3742.
- [12] R. Gayle, P. Segars, M. C. Lin, and D. Manocha, “Path planning for deformable robots in complex environments,” in *Proc. Robotics: Science and Systems (RSS)*, 2005, pp. 225–232.
- [13] K. G. Gopalakrishnan and K. Goldberg, “D-space and deform closure grasps of deformable parts,” in *International Journal of Robotics Research*, vol. 24, no. 11, 2005, pp. 899–910.
- [14] S. Gottschalk, M. C. Lin, and D. Manocha, “OBB-tree: A hierarchical structure for rapid interference detection,” 1996. [Online]. Available: <http://www.cs.unc.edu/~geom/SSV>
- [15] J. W. Hill, P. S. Green, J. F. Jensen, Y. Gorfou, and A. S. Shah, “Telepresence surgery demonstration system,” in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 1994, pp. 2302–2307.




Fig. 5. Scenario 1: The objective is to part the tissue flaps on either side of the incision and completely expose the underlying area of interest (red). Optimal initial grasp locations as determined by our method (a). Figures (b)–(d) illustrate the trajectory of the two grippers as the tissue is retracted while minimizing the maximum deformation energy and avoiding the obstacle in the workspace.




Fig. 6. Scenario 2: Grasp location and gripper trajectory that minimizes control effort. The gripper starts on the other side of the obstacle. (a) The gripper performs a twist to minimize displacement, thereby inducing large deformations. (b) The gripper rectifies itself and performs the retraction successfully (c–d).




Fig. 7. Scenario 2: Grasp location and gripper trajectory that minimizes the maximum tissue stress. The gripper starts on the side of the obstacle (a). The gripper maneuvers around the obstacle and raises the tissue flap straight up in order to minimize torsional stress in the constrained nodes (b–d).

- [16] S. Hirai, T. Tsuboi, and T. Wada, “Robust grasping and manipulation of deformable objects,” in *Proc. IEEE International Symposium Assembly and Task Planning*, 2001.
- [17] A. M. Howard and G. A. Bekey, “Intelligent learning for deformable object manipulation,” in *Autonomous Robots*, vol. 9, 2000, pp. 51–58.
- [18] R. Jansen, K. Hauser, N. Chentanez, F. van der Stappen, and K. Goldberg, “Surgical retraction of non-uniform deformable layers of tissue: 2D robot grasping and path planning,” in *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*, 2009, pp. 4092–4097.
- [19] F. Lamiroux and L. E. Kavraki, “Planning paths for elastic objects under manipulation constraints,” in *International Journal of Robotics Research*, vol. 20, no. 3, 2001, pp. 188–208.
- [20] S. M. LaValle, *Planning Algorithms*. Cambridge, U.K.: Cambridge University Press, 2006, Available at <http://planning.cs.uiuc.edu>.
- [21] M. Moll and L. E. Kavraki, “Path planning for deformable linear objects,” in *IEEE Transactions on Robotics*, vol. 22, no. 4, 2006, pp. 625–636.
- [22] A. Nealen, M. Müller, R. Keiser, E. Boxerman, and M. Carlson, “Physically based deformable models in computer graphics,” in *Computer Graphics Forum*, vol. 25, no. 4, 2006, pp. 809–836.
- [23] H.-W. Nienhuys, “Cutting in deformable models,” Ph.D. dissertation, Utrecht University, 2003.
- [24] G. Picinbono, H. Delingette, and N. Ayache, “Nonlinear and anisotropic elastic soft tissue models for medical simulation,” in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2001, pp. 1370–1375.
- [25] S. Rodriguez, J.-M. Lien, and N. M. Amato, “Planning motion in completely deformable environments,” in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2006, pp. 2466–2471.
- [26] M. Saha and P. Isto, “Manipulation planning for deformable linear objects,” in *IEEE Transactions on Robotics*, vol. 23, no. 6, 2007, pp. 1141–1150.
- [27] J. R. Shewchuk, “An introduction to the conjugate gradient method without the agonizing pain,” Carnegie Mellon University, Tech. Rep., 1994. [Online]. Available: <http://portal.acm.org/citation.cfm?id=865018>
- [28] H. Si, “Tetgen,” 2007. [Online]. Available: <http://tetgen.berlios.de/>
- [29] R. H. Taylor, *et al.*, “An image-guided robotics system for precise orthopaedic surgery,” in *IEEE Transactions on Robotics and Automation*, vol. 10, no. 3, 1994, pp. 261–275.
- [30] O. C. Zienkiewicz, R. L. Taylor, and J. Z. Zhu, *The Finite Element Method: Its Basis and Fundamentals*. New York: Butterworth-Heinemann, 2005.