

Automated Tissue Retraction for Robot-Assisted Surgical Procedures

Sachin Patil

Ron Alterovitz

Abstract—Robotic surgical assistants (RSAs) are enhancing physician performance, enabling them to perform more delicate and precise minimally invasive surgery (MIS). However, these devices are currently tele-operated and lack autonomy. In this paper, we focus on automating a relatively simple but commonly performed surgical task of *tissue retraction*, which involves grasping and lifting a thin layer of tissue to expose an underlying area. Specifically, we target the problem of computing an optimal plan and an optimal grasp location for automated retraction of tissue flaps using a generic 6-DOF gripper. We use a nonlinear FEM (Finite Element Method) formulation 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 a globally optimal plan. We propose novel cost metrics which are relevant for evaluating the optimality of the plans in a clinical setting, including minimizing the maximum deformation energy, minimizing the maximum tensile stress, and minimizing the control effort in lifting the tissue flap. We present experimental results on challenging retraction scenarios that include obstacle avoidance and heterogeneous tissues.

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 (MIS). Studies have shown that these robots can improve procedure success rates, reduce bleeding, and decrease recovery time [9]. As a result, these robots are increasingly being adopted in hospitals, 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 require direct control by physicians, i.e. they lack autonomy.

Integrating motion planning with RSAs has the potential to enable RSAs to perform certain motions autonomously. 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: manipulate the tissue to provide sufficient exposure and, at the same time, avoid excessively large forces that damage the tissues.

Enabling RSAs to autonomously perform sub-procedures like tissue retraction could have a significant impact on patient care. Motion planning can provide a form of “auto-pilot” for




Fig. 1: *Tissue retraction procedure: (a) Given an initial grasp location and a camera viewpoint, the tissue flap is manipulated by the gripper such that the area of interest (marked by the red circle) is completely visible (b) Top view of the scene indicating the configuration of an obstacle in the workspace (which protrudes from the side and lies above the slab of tissue).*

medical robotic devices, enabling them to complete simple, routine tasks automatically. This 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. Autonomous control would also enable RSAs to simultaneously control more than the two manipulators that can be directly controlled by human hands, enabling greater dexterity and faster procedure completion times. The ability to perform motions autonomously would also facilitate remote surgery in remote or hostile environments in situations where time lag due to network limitations prohibits effective tele-operation.

We present an approach to compute plans for automated tissue retraction using a generic 6-DOF gripper. We consider three novel optimization objectives relevant to medical applications: minimizing the maximum deformation energy, minimizing maximum tissue stress and minimizing the cumulative control effort. The method uses a nonlinear finite element method (FEM) to compute equilibrium states of the tissue flap corresponding to sampled gripper configurations. These states are used to construct a global roadmap to explore the space of deformations and compute an optimum plan that avoids obstacles. A fast visibility test using occlusion queries on graphics processing units (GPUs) is used to compute the proportion of visible area beneath the tissue flap for determining goal states. Our planner computes trajectories for the gripper to accomplish the tissue retraction task while avoiding obstacles, as shown in Fig. 1.

II. RELATED WORK

Robotic Surgical Assistants (RSAs): In the context of medical robotics, tele-robotic or master/slave systems (such as Intuitive Surgical’s da Vinci® system [8]) operate under direct

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

control of the surgeon, allowing for greater precision and dexterity in difficult to operate environments [8], [14]. A few fully 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 [26]. Our approach proposes to autonomously perform the task of tissue retraction for deformable tissues.

Modeling of Deformable Objects: Physically-based simulation of deformable objects is a well-studied area in solid mechanics [27] and computer graphics [21]. The choice of the underlying model is application-specific and has a substantial influence on the physical accuracy of the estimated deformations. Lumped mass-spring systems are easy to implement and can be simulated efficiently but there is no intuitive relation between the spring constants and physical material properties [21]. Finite Element Methods (FEM) are based on the theory of continuum mechanics [27] and simulate deformations accurately at the cost of increased computational complexity. 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 [23], [22].

Motion Planning and Manipulation of Deformable Objects:

Robot motion planning and manipulation are active areas of research and have historically focused on robot motion planning and manipulation for robots composed of rigid links operating in environments with rigid objects [19]. Recent work has begun to explore motion planning for deformable robots in static environments. Lamiroux et al. [18] 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. [5] 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 planning scheme for deformable objects modeled as mass-spring systems. Rodriguez et al. [24] 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 [20] and flexible needle devices traveling through deformable tissue [4]. Prior work on robot grasping and manipulation have generally assumed that the grasped object is rigid [6]. Howard et al. [16] model deformable parts using a mass-spring model and use a neural network to control the gripper. Hirai et al. [15] propose a robust control law for local manipulations of deformable parts using tactile and video feedback. 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. [17] propose a geometric approach to compute candidate grasp locations and optimal trajectories for automated tissue retraction. This computationally fast approach considers a 2D cross-sectional

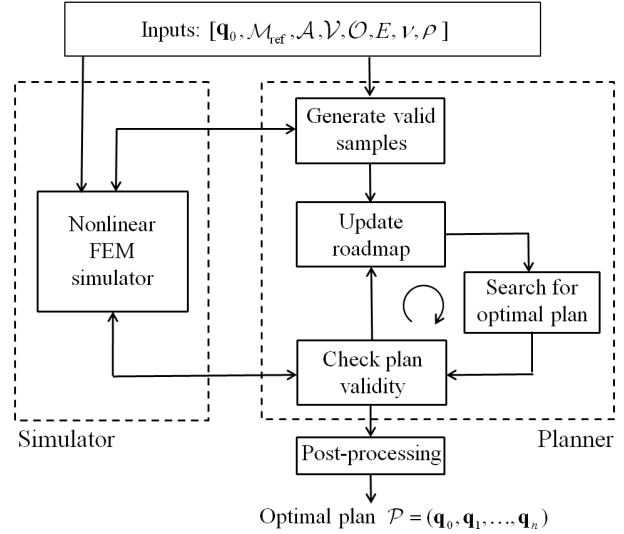


Fig. 2: System Overview: The PRM-based global planner (Planner) relies on the deformable body simulator (Simulator) to compute static, equilibrium states of the tissue flap corresponding to randomly sampled gripper configurations and generate an initial roadmap. The roadmap is queried for an optimal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

model of the tissue flap and computes trajectories for the gripper in a plane. In contrast, our approach works with a fully 3D tissue flap model and handles generic contact and manipulation constraints, and searches for optimal solutions in constrained environments containing obstacles that need to be avoided while performing the procedure.

III. PROBLEM STATEMENT

We assume that the geometry of the tissue flap 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). For simplicity, it is assumed 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 the 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 manipulating the tissue flap. We assume that the tissue moves slowly as it is manipulated by the gripper and this process can be approximated as a quasi-static process and dynamics is currently ignored. We also assume that the motion of the gripper and the subsequent deformations of the tissue flap are completely deterministic; we plan to consider the effects of uncertainty due to factors such as material parameters, actuation or slippage in future work.

It is important to address the notion of optimality for tissue retraction since exerting excessive forces at the gripper and large deformations of soft tissue can cause irreversible damage such as tearing of tissue. We propose the following cost metrics for computing optimal plans:

- **Minimizing the 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 other cause irreversible tissue damage. Paths with lower deformation energies will be less susceptible to these effects.
- **Minimizing the maximum tensile stress:** The maximum tensile stress encountered in the deformed body represents an upper limit to the forces that can be applied to deformable objects without causing phenomena such as tearing and fracture (also known as the yield point) [27]. Minimizing this quantity helps discard states that are very close or beyond the yield points, resulting in *safe* plans.
- **Minimizing the total control effort:** The effort expended in performing the task, which we define as the total control effort, is a combination of the total deformation energy and the gripper displacement (measure of mechanical effort). This metric balances the search between two competing objectives: minimizing the deformation energy of the intermediate states along the path and minimizing the total mechanical effort in executing the plan.

The choice of the cost metric depends on several factors such as maneuverability of the gripper in the workspace and type of tissue involved in the retraction task, and is determined by the surgeon during the procedure.

The problem can now be formally stated as follows:

Objective 1: To determine a feasible plan in terms of the control inputs (configurations encoding both the position and orientation of the gripper), such that the area of interest beneath the tissue flap is completely visible, while minimizing the chosen cost metric.

Input: Initial configuration of the gripper \mathbf{q}_0 , tissue flap 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 (Young’s modulus E , Poisson’s ratio ν and density ρ).

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

Objective 2: The above framework (Fig. 2) assumes that the initial grasp location of the gripper is known. It can be generalized to also compute the optimal grasp location for the gripper $\mathbf{q}_0^{\text{opt}}$, which minimizes the cost metric over all possible grasp locations for a given scenario.

IV. MANIPULATION AND PATH PLANNING

The two key components of our approach are the sampling-based global planner (Planner) and the nonlinear finite element simulator (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 (Sec. V). These gripper configurations, coupled with corresponding equilibrium mesh states, are used to explore the space of deformations and compute an *optimal* plan (Sec. IV). 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 choice of the simulator.

A. Generating Valid Samples

Manipulation Constraints: The tissue flap 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.

Manipulation constraints of the gripper are generated by randomly sampling configurations $\mathbf{q}_i \in SE(3)$ of the 6-DOF gripper. All the nodes that lie within the bounding box of the gripper jaws are fixed. For a given sample \mathbf{q}_i , their displacement is computed by applying the transformation given by \mathbf{q}_i to their positions in the reference state \mathcal{M}_{ref} . An equilibrium mesh state \mathcal{M}_i corresponding to the gripper configuration \mathbf{q}_i can then be computed (Sec. V).

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 [2] to discretize the mesh and reliably detect self-penetration.

Sample Optimization: Since the gripper configuration \mathbf{q}_i is sampled randomly, high stresses might build up in the elements constrained by the gripper (as shown in Fig. 3(a)), resulting in undesirable high-stress states. For every valid sample generated, we perform a local gradient-descent based optimization on the rotation 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(b)). Our experiments indicate that fewer states are eventually required in the roadmap to compute a feasible plan if this local improvement is applied.




Fig. 3: *Local optimization: (a) High stresses build up in the elements constrained by the gripper corresponding to a randomly generated configuration (b) Local gradient-descent based optimization is performed on the gripper rotation to minimize the stresses in the constrained elements.*

User-Guided Heuristic Sampling: The pre-processing time involved in computing equilibrium mesh states (corresponding to randomly sampled gripper configurations) is prohibitive. A user-guided heuristic sampling scheme is adopted to reduce this computation time. The user (surgeon in this case) indicates a probable retraction path for the gripper. In our experiments, the user indicates a final gripper position (in \mathbb{R}^3) and the trajectory is specified as a straight line joining the initial and user-specified gripper position. The user can also choose to specify the trajectory in terms of multiple way-points. The samples \mathbf{q}_i are generated in a user-defined neighborhood of this trajectory. It should be noted that the quality of the final solution obtained will depend on the user-specified trajectory and the choice of neighborhood distribution.

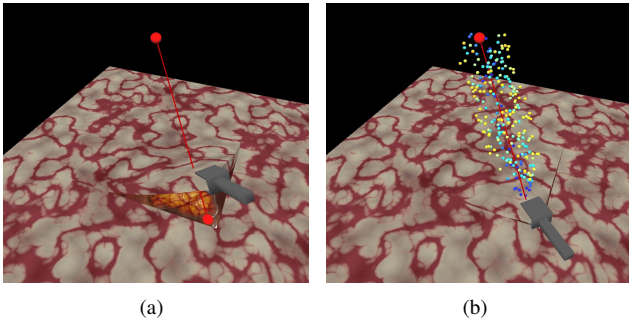


Fig. 4: User-guided heuristic sampling: (a) User (surgeon) specifies an estimate of the final position of the gripper and a simple trajectory is obtained by connecting the initial and user-specified gripper position (in \mathbb{R}^3) by a straight line. (b) Gripper configurations \mathbf{q}_i are sampled within a user-defined neighborhood of the trajectory.

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} and 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. They are exact but computationally expensive. Image-space methods compute the area visible in terms of pixels, when the entire scene is rendered as an image from the camera viewpoint. They are approximate but fast, and there is no need to determine the camera parameters in advance. The pixel accuracy suffices since we are only interested in the exposure of a pre-defined area of interest. We implemented the visibility computation using occlusion queries on graphics processing units (GPUs) [1]. 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 set of goal configurations \mathbf{Q}_g . We

also store the mesh state \mathcal{M}_i along with each sample \mathbf{q}_i for the purposes of computing the distance metric defined below.

Distance Measure: We use a heuristic distance measure [18] 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 [19]. 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: A path can be defined between two samples represented by configurations \mathbf{q}_i and \mathbf{q}_j by linearly interpolating between the two configurations in $SE(3)$ [19]. The edge connecting any two samples in the roadmap is considered to be *valid* if each of these intermediate mesh states is *valid* (Sec. IV-A) and if both the mesh and the gripper do not collide with any obstacles in the workspace. We use PQP [13] for collision detection with obstacles.

D. Assigning Roadmap Edge Weights

The cost metrics outlined in Sec. III are used for assigning weights to edges in the roadmap and 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 $\eta \in \mathcal{N}_i$. The total deformation energy Π_i of the mesh state \mathcal{M}_i is: $\Pi_i = \sum_{\tau \in \mathcal{M}_i} v(\tau) \cdot W(\mathbf{E}) + \sum_{\eta \in \mathcal{N}_i} \mathbf{f}_{\text{ext}}^{\eta} \cdot \mathbf{u}_{\eta}$, where $v(\tau)$ is volume of the tetrahedron τ , \mathbf{u}_{η} is displacement of node η , \mathbf{E} is the Green strain tensor, $W(\mathbf{E})$ is the energy density function for the chosen material model [27], [22].

Minimizing maximum tissue stress: The cost of a roadmap edge is chosen to be the maximum stress 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 one 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 $\max(\text{abs}(\sigma^i)), i \in \{1, 2, 3\}$ and the yield condition (from maximum principal stress theory) [27] states that yield (tearing in this case) occurs when this value is greater than the tensile yield strength of the material.

Minimizing control effort in lifting tissue flap: We express the total control effort C using the metric defined in Moll et al. [20] along a specified path $(\mathbf{q}_1, \dots, \mathbf{q}_n)$ as: $C = \sum_{i=1}^{n-1} d(\mathbf{q}_i, \mathbf{q}_{i+1})(\Pi_i + \Pi_{i+1})/2$. 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 planning requires running the simulation along the interpolated path and is a computationally expensive process. To reduce computation times, we use a *lazy* local planning scheme [7]. An initial estimate of the roadmap (constructed by assigning estimates of edge weights as described) is searched for a minimum-cost path from \mathbf{q}_s to \mathbf{Q}_g . The local planner is then invoked to check if all the edges along the path are *valid*. Any *invalid* edges are discarded and the *valid* edges are marked as processed and their edge weights updated. This process is repeated till a valid (optimal) plan is obtained or no solution is found. Since our method relies on probabilistic roadmaps to search for a valid solution, it is probabilistically complete. As the number of samples generated in the pre-processing phase increase, the solution converges to the true optimal.

We use the Dijkstra’s shortest path algorithm for searching for an optimal path from \mathbf{q}_0 to \mathbf{Q}_g for additive metrics (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 Bottleneck Shortest Paths problem [3] and can be solved by a minor modification to the Dijkstra’s shortest path algorithm.

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 \mathbf{q}_i using a cubic spline [10] to maintain C^1 -continuity between configurations.

V. SIMULATION OF TISSUE FLAP RETRACTION

Simulating soft tissue is challenging since it must provide convincingly realistic visual and haptic response to manipulations such as grasping, pulling and cutting. We chose to model tissue deformations using nonlinear continuum mechanics (FEM) to account for large tissue deformations. The tissue itself is assumed to be isotropic, nearly incompressible and exhibits hyper-elasticity according to the St. Venant Kirchhoff material model (commonly chosen to represent bio-mechanical deformable objects [23], [22]). We refer the reader to [27] 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 [23]. Contact with the under-

lying base (modeled as a plane in our experiments) is handled by a penalty-based scheme where reaction forces proportional to the penetration depth of the intersecting nodes, attempt to resolve the collision. More sophisticated collision resolution schemes could also be used [11]. We do not consider the interaction forces and stick-slip friction between the gripper and the tissue flap.

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 [22] 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 [23]. This allows us to model the nearly incompressible nature of soft tissue and to prevent inversion of tetrahedra under strong constraints.

Solving for Static Equilibrium: The tissue flap is assumed to be anchored at a subset of nodes (usually along one or more edges along the boundary of the mesh) and is also manipulated by the gripper. Let \mathcal{N}_{fix} denote the set of constrained nodes that have a fixed displacement. This set comprises of both *anchored* nodes (which have zero displacement) and *grasped* nodes (the displacement of which 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 now given by solving the (reduced) system of equations, $\Phi(\tilde{\mathbf{u}}) = \tilde{\mathbf{f}}_{\text{ext}}$, where $\tilde{\mathbf{u}}$ now 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 gradients solver [25].

VI. RESULTS

The planner was implemented in C++ and tested on a 3.2 GHz 8-core Intel® i7™ workstation.

We use a rectangular tissue flap model of dimensions 5 cm \times 5 cm \times 0.25 cm for our experiments. The tissue flap is discretized into linear, tetrahedral elements and the resulting mesh contains 1300 nodes and 4000 elements. The tissue density ρ is set to 1000 kg/m³, Poisson’s ratio ν to 0.45 and Young’s Modulus E is set to 1 Mpa (unless otherwise specified). Gravity acting on the tissue flap is set to 9.8 m/s² acting downwards. We tested our method on two tissue retraction scenarios:

Scenario 1: Fig. 5 shows an incision in the tissue from an overhead viewpoint. The tissue is assumed to be *homogeneous* throughout with a Young’s Modulus (E) of 1 MPa. The objective here is to part the two tissue flaps to expose an arbitrarily shaped area of interest while minimizing the maximum deformation energy and avoiding the obstacle. Since the two tissue flaps are symmetric in the initial boundary conditions (each tissue flap is anchored along three edges), equilibrium mesh states are computed in the pre-processing stage for a single tissue flap. The optimal initial grasp configuration $\mathbf{q}_0^{\text{opt}}$

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 complex scenario in which the tissue flap is anchored along two of the edges and the gripper is allowed to grasp the flap anywhere along the other two free edges. An obstacle protrudes from the side and extends over the tissue flap. The tissue flap is modeled as a *heterogeneous* structure with veins running through its entire length. The veins 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 veins to be 1 MPa and for the surrounding tissue to 250 KPa. The optimal initial grasp configuration $\mathbf{q}_0^{\text{opt}}$ is obtained by performing an exhaustive search over a set of 9 possible initial grasp configurations. Figures 6 and 7 show the optimal grasp locations and the retraction trajectories for the gripper when the cost metric aims to minimize the maximum stress and the control effort respectively. In both cases, the gripper must pull the tissue to avoid obstacle intersection.

Scenario	Num samples	Sample generation (hrs)	Lazy planning			Total (secs)
			Collision detection (secs)	Mesh validation (secs)	FEM sim (secs)	
1	1000	1.212 (\pm 0.149)	0.563	24.94	7.113	32.713 (\pm 5.98)
2	1000	3.027 (\pm 0.225)	1.803	32.85	48.147	83.316 (\pm 38.98)

TABLE I: Performance of our approach. Provided times are averaged over multiple initial grasp locations $\mathbf{q}_0^{(i)}$. Standard deviation times are provided in brackets. Times reported are for parallel sample generation (using OpenMP) and serial (lazy) planning.

Scenario	Random sampling		User-guided sampling		Metric RMS error	
	Num samples	Sample generation (hrs)	Num samples	Sample generation (hrs)	Min. maximum energy	Min. maximum stress
1	1000	1.212 (\pm 0.149)	200	XXX (\pm YYY)	XXX	XXX
2	1000	3.027 (\pm 0.225)	200	XXX (\pm YYY)	XXX	XXX

TABLE II: Comparison of random sampling and user-guided sampling schemes. The last column reports the RMS error in the chosen cost metric for the user-guided sampling scheme (as compared to the random sampling scheme).

VII. CONCLUSION

We have described a framework for automating the task of tissue retraction in robot-assisted surgical procedures. Given a model of the tissue flap, a nonlinear FEM 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 three metrics: minimizing the maximum deformation energy, minimizing maximum stress, and minimizing the control effort

in lifting the tissue flap. The method is directly applicable to arbitrary shaped tissue flaps.

In future work, we plan to improve the speed of the method by improving the simulation. Recent research in deformable body simulation (such as modal analysis [27]) can be used to achieve significant speedups. Another important issue that needs to be addressed is the discrepancy in the results obtained and observed tissue behavior due to uncertainty in the chosen material parameters, gripper actuation or tool-tissue interaction. A global roadmap provides an excellent starting point in terms of incorporating this uncertainty (similar to the Stochastic Motion Roadmap framework suggested in [4]). Empirical data obtained from tool-tissue interaction studies can be incorporated for modeling realistic gripper forces.

VIII. ACKNOWLEDGEMENTS

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

REFERENCES

- [1] “ARB occlusion query,” 2002. [Online]. Available: http://www.opengl.org/registry/specs/NV/occlusion_query.txt
- [2] “Tetgen,” 2007. [Online]. Available: <http://tetgen.berlios.de/>
- [3] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms and Applications*. Prentice Hall, 1993.
- [4] R. Alterovitz and K. Goldberg, *Motion Planning in Medicine: Optimization and Simulation Algorithms for Image-Guided Procedures*. Springer, July 2008, vol. 50.
- [5] 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.
- [6] A. Bicchi and V. Kumar, “Robotic grasping and contact: A review,” in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2000, pp. 348–353.
- [7] R. Bohlin and L. E. Kavraki, “Path planning using lazy prm,” in *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2000, pp. 521–528.
- [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] 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>
- [14] 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.
- [15] S. Hirai, T. Tsuboi, and T. Wada, “Robust grasping and manipulation of deformable objects,” in *Proc. IEEE International Symposium Assembly and Task Planning*, 2001.
- [16] A. M. Howard and G. A. Bekey, “Intelligent learning for deformable object manipulation,” in *Autonomous Robots*, vol. 9, 2000, pp. 51–58.
- [17] 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.
- [18] F. Lamiraux 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.




Fig. 5: *Scenario 1: The objective is to part the two tissue layers involved in the incision and completely expose the arbitrarily shaped region of interest underneath. (a) Optimal initial grasp locations as determined by our method. 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: Optimal grasp configuration and trajectory of the gripper that minimizes the maximum tissue stress. (a) The gripper starts on the same side as that of the obstacle. (b)-(d) Notice how the gripper manoeuvres around the obstacle and raises the tissue flap straight up in order to minimize any torsional stress in the constrained nodes.*




Fig. 7: *Scenario 2: Optimal grasp configuration and trajectory of the gripper that minimizes the control effort: (a) The gripper starts off on the other side of the obstacle. (b) In a bid to minimize the gripper displacement, the gripper performs a twist inducing a high-energy state. (c) - (d) The gripper rectifies itself and performs the retraction successfully.*

- [19] S. M. LaValle, *Planning Algorithms*. Cambridge, U.K.: Cambridge University Press, 2006, Available at <http://planning.cs.uiuc.edu>.
- [20] 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.
- [21] 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.
- [22] H.-W. Nienhuys, "Cutting in deformable models," Ph.D. dissertation, Utrecht University, 2003.
- [23] 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.
- [24] 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.
- [25] 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>
- [26] R. H. Taylor, B. D. Mittelstadt, H. A. Paul, W. Hanson, and P. K. 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.
- [27] O. C. Zienkiewicz, R. L. Taylor, and J. Z. Zhu, *The Finite Element Method: Its Basis and Fundamentals*. New York: Butterworth-Heinemann, 2005.