

Chapter 5

A Delaunay approach to interactive cutting in triangulated surfaces

5.1 Introduction

Our approach to deformation is based on the Finite Element Method (FEM). In this method, mesh size determines the computational requirements of a simulation. Larger meshes result in more degrees of freedom in the discretized problem, so solutions take more time to compute. This has motivated our work in Chapter 3, where we have tried a technique for simulating cuts that does not increase mesh size like subdivision techniques do. This technique creates flat elements in the mesh as artefacts, and we have found that they cause a considerable slowdown of the Linear CG algorithm. In Subsection 4.5.6 we have seen that element shape also affects the nonlinear CG and dynamic relaxation algorithms. Hence, mesh change operations, such as simulated cuts or cauterizations, should not only keep mesh size low; they should also keep elements of the mesh well-shaped. In this chapter¹ we address this problem for cutting in triangulations, by presenting a method that keeps mesh size low and keeps mesh quality high.

A FEM discretization is a form of interpolation: the continuous unknown in the problem is interpolated, so that the differential equation is transformed into a set of equations for a finite number of variables. The original problem is a differential equation, hence it is important that the derivative of the solution is approximated well by the interpolation. When we look at the influence of element shape on the derivative, we see that the large angles cause unbounded errors in the derivative. This is illustrated in 2D in Figure 5.1, and a similar argument also holds in 3D. Therefore, large angles should always be avoided. The convergence speed of iterative algorithms for linear problems

¹This chapter was based on a paper published at the Fifth International Workshop on Algorithmic Foundations of Robotics (WAFR 2002) [74].

is related to the condition number of the stiffness matrix, as shown in Sections 4.1 and 2.4. The condition number is the ratio of largest and smallest eigenvalue of the stiffness matrix, and these eigenvalues can be bounded by eigenvalues of separate elements: each element can also be seen as a separate elastic object, whose deformations are also described by linear elasticity, and their elastic behavior can be condensed in a small *element stiffness matrix*. Therefore, mesh quality can be optimized by optimizing the shape of individual elements. In a second-order elliptic partial differential equation (like linear elasticity), both very large and very small angles, and small elements cause high condition numbers [87]. Therefore, a mesh with ‘round’ elements, i.e. no angles close to π and 0, and uniform element sizes is a good general purpose mesh.




Figure 5.1: The derivative of a linear function with values $f(a)$, $f(b)$ and $f(c)$ in triangle abc in the plane has directional derivative $\partial f/\partial y = \frac{f(c) - (f(a) + f(b))/2}{H}$. When $H \rightarrow 0$, the derivative tends to infinity.

Simulation of cuts in surgery simulation is related to simulation of other destructive surgical procedures. The first operation to have been simulated on volumetric meshes is cauterization. This was done by removing elements in contact with a virtual cauterization tool [26]. A disadvantage of element removal is that it produces a jagged surface on the virtual tissue.

For cutting, subdivision methods are the norm [11, 46, 66]: elements that are in contact with the scalpel are marked active, and subdivided to produce a cut conforming to the scalpel position. The subdivision moves along with the scalpel during its stay in an active tetrahedron. When the scalpel leaves an active element, the subdivision is entered in the mesh permanently. A 2D example of subdivision is in Figure 5.2. Subdivision methods always increase the size of the mesh. Moreover, these methods tend to produce degeneracies. This is caused by the use of an *active region*. Mesh modification is done only within a fixed region of the mesh, and if the scalpel moves close to the boundary of that region, poorly-shaped elements are inevitable. Ganovelli and O’Sullivan [45] try to counter the degeneracies caused by subdivision cutting. They deal with these degeneracies by collapsing short edges of the mesh. This approach does improve the quality of the mesh, but this solution does not repair all inconsistencies: not all edges may be contracted, and flat triangles and tetrahedrons, which do not contain short edges but are still degenerate, are not dealt with.

In Chapter 3 we have tried an approach where scalpel nodes are snapped to the trajectory swept by the scalpel. The advantage of this method is that the mesh size remains small, and few short edges are created. However, there are a number of disadvantages: since no new nodes are created, the resolution of the cut is bounded by the mesh resolution. The incision does not reach up to the position of the scalpel, but lags behind it. A more serious problem is that snapping can result in degenerate elements in the

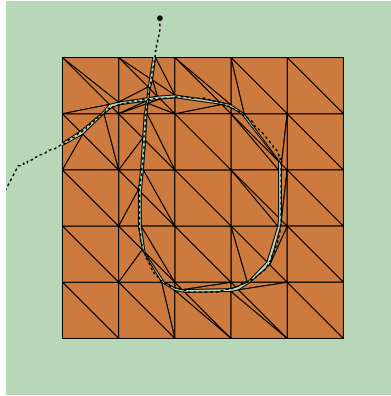


Figure 5.2: A cut produced by a subdivision method. The cut produces many small and flat triangles.

mesh. Such degeneracies are dealt with by subdividing flat elements, and collapsing the resulting short edges, effectively removing the flat element. Unfortunately, not all edges can be collapsed. Using existing mesh features as a basis for mesh modification is problematic: when the scalpel is not especially close to a mesh feature, it may not be possible to match the mesh topology to the scalpel path without introducing degeneracies.

In summary, producing high quality cuts in tetrahedral meshes is a difficult problem. For cutting in surfaces, there is an analogous problem, which has the same difficulties as volumetric cutting approaches. Serby et al. [84] propose a method which also relies on a form of node snapping: the scalpel is modeled as a line segment, and nodes from the mesh are projected onto that segment. In a post-processing step, the edge lengths and element volumes are optimized using a particle system. The result is a good looking cut, without a decrease in mesh quality and size. However, in reality, the path of a scalpel is not a large line segment, but a concatenation of several small ones, and their approach does not seem to address this issue. Bruyns et al. [21] use surface cutting with subdivision in large-scale simulations of various procedures.

In light of the complications of volumetric cutting, we will take a step back, and analyze the cutting problem for surface meshes first. In this chapter we will present a method that produces cuts in a triangulated surface which does not decrease the mesh quality and keeps the mesh size low. We have analyzed this problem primarily to gain insight into the cutting problem for 3D tetrahedral meshes. Nevertheless, this technique could be applied in surgery simulation for membrane-like structures, such as skin or intestine.

5.2 Cutting in 2D

We state the general mesh cutting problem as follows: given a starting mesh, and positions of a user-controlled scalpel, modify the mesh at every moment to show an incision that represents the past trajectory of the scalpel, and ends exactly at the scalpel. The

challenge in this problem is to produce a well-shaped mesh with few elements. The simplest case is the two-dimensional form. Here the mesh is a triangulation in the plane, and the virtual scalpel is a point. The challenge is to produce elements that have no large angles and no short edges.

5.2.1 Delaunay triangulation

Since we are putting emphasis on the quality of the mesh, we briefly review the Delaunay Triangulation (DT), a popular technique for generating a well-shaped triangulation of a given set of points. The DT of a set of points is defined as a triangulation where the circumcircle of every triangle does not contain any other points from the set. This property is also referred to as the empty-circle property. The empty-circle property can also be defined for edges of the mesh: an edge is called *legal* or *Delaunay* when the circumcircles of its incident triangles do not contain the opposite node of the other triangle. An example of an illegal edge is in Figure 5.3. Delaunay triangulations and Delaunay edges are intimately related: a Delaunay triangulation only has Delaunay edges.

Non-Delaunay or illegal edges of a triangulation can always be flipped: the two triangles incident with the edge always form a convex quadrilateral, and the diagonal may be switched. The flipped diagonal is always Delaunay, and the minimum angle of the pair of triangles is always increased, thus improving mesh quality. Flipping illegal edges only affects a single part of the mesh, so it can be seen as local improvement strategy. The Delaunay triangulation can be constructed by starting with an arbitrary triangulation and flipping illegal edges until none are left. The final result maximizes the minimum element angle. This is called the maxmin-angle property.

There are various ways to measure element quality: for example, minimum angle, circumradius to shortest edge ratio, circumradius to inscribed radius ratio, etc. In 2D all these measures are equivalent up to constant factors [63], so the maxmin-angle property of the DT means that it is a reasonable meshing technique for virtually all element quality measures. When coupled with algorithms for point insertion it is a basis for many refinement meshing techniques [10, 25, 82].




Figure 5.3: An illegal edge has a triangle whose circumcircle contains the opposite vertex of a neighboring triangle (left). By reconnecting (‘flipping’) the edge, the new circumcircles only contain the vertices it circumscribes (right).

5.2.2 Cutting and Delaunay flips

Our problem, cutting in meshes, is different from the standard meshing problem. A starting mesh is given, and the shape of our domain is variable: as the cut progresses, the shape of the boundary is changed. The Delaunay Triangulation assumes a given set of points, and it is always convex. Cuts result in non-convex boundaries, so the standard DT is not directly usable. Nevertheless, we can retain the idea of using edge flips to locally improve the mesh.

Our approach works as follows: when cutting, the scalpel is attached to a node, the *active node*, so moving the scalpel moves the active node. The active node is always part of the boundary, so it is incident to two boundary edges, the *cut edges*. During a cut, these edges almost coincide geometrically, and form an *incision*. We call triangles incident with the active node *active triangles*. The active node is moved, and after each movement, local remeshing is applied to the active triangles. The remeshing process consists of the following actions.

- Edges are flipped to improve the element angles.
- If an internal node is found close to the active node, it is removed.
- The incision is split, thus introducing new nodes to approximate the scalpel path.

The net effect of the last two actions is that nodes are removed in front of the scalpel, and inserted behind the scalpel. The technique is demonstrated in Figure 5.4; a more elaborate example is in Figure 5.8.

The Delaunay criterion tends to flip away triangles with angles approaching 180° , since these have large circumcircles. Nodes that are close together get connected. Such triangles have a small angle opposite a short edge. These triangles are also undesirable, so they are removed by by the second action.

When the scalpel enters the mesh, it is close to the boundary of the mesh, so a realistic incision would contain very short edges. To prevent these short edges, creation of incisions is postponed. When the scalpel enters the mesh close to an existing node, the node is moved and labeled active. If it enters in somewhere else, a new node marked active is inserted at the entry point. The active node moves along with the scalpel, creating a temporary dent. When the active node is sufficiently far from the entry point, it is “fixed up”: nodes are added at the entry point, creating an incision. The procedure is shown in Figure 5.6. If the scalpel is *retracting* (moving away from the object) before this fix-up happens, the dent is left in the mesh permanently.

After the entry fix-up, the cut edges are in almost in the same location. From a geometric point of view, we can identify both cut edges into a single edge, and check if this edge satisfies the empty-circle criterion. If it does not, new nodes are inserted, where the line connecting nodes opposite the cut edges intersect the scalpel path. We call this procedure a *incision split*, and it is demonstrated in Figure 5.5. When the scalpel passed that point previously, close nodes were removed, so the split will not lead to short edges. The newly inserted nodes are dilated slightly, to prevent numerical problems when a path self-intersects. The new nodes always lie on a line that connects two existing mesh nodes, hence the accuracy of the represented trajectory is bounded by the resolution of the starting mesh.




Figure 5.4: The evolution of a simple 2D cut is depicted from left to right. The key steps are to flip edges (center), to remove close nodes (left), and to insert nodes behind the scalpel (right).




Figure 5.5: During an incision split, nodes are inserted in the cut edges. This happens when ac , with $c = (c_1 + c_2)/2$, fails the empty-circle criterion as diagonal of aco_1o_2 . The newly inserted nodes c'_1 and c'_2 are inserted where o_1o_2 intersects the scalpel path.




Figure 5.6: Incisions are postponed when the scalpel enters the mesh.

A similar situation arises when the scalpel exits. In this case, the scalpel comes arbitrarily close to the boundary of the mesh. To prevent arbitrarily flat elements from occurring, the exit point is predicted, and the cut is finished before the scalpel actually hits the boundary. At every step, the movement of the scalpel is extrapolated. When this extrapolation hits the boundary, and is close to the active node, the cut is finished: a node is inserted at the extrapolation, and the cut is dissected. This exit procedure still leaves relatively short edges permanently in the mesh. To rectify this, these edges are contracted. The procedure is shown in Figure 5.7 The boundary edges that are created with this exit operation are added to a list of forbidden edges. These edges are not tested for collisions during ensuing movements. This prevents artifacts when the actual scalpel movement differs from the predicted movement. If the scalpel is sufficiently far from the exit point, they are eligible for collision checking again.

5.3 Surface cuts in 3D

Triangulated surfaces in 3D are a common tool in computer graphics. They have also been used for surgery simulations [19, 22]. In these cases elastic behavior was simulated with damped mass-spring systems instead of the FEM. Nevertheless, the concerns for mesh quality continue to hold: flat elements are inverted more easily, and small elements correspond to short springs with small masses. Their high vibration frequencies translate into small time steps, which makes time integration expensive.

In the 2D scheme, the scalpel is a point, and it is attached to a single active node. Triangles are remeshed in the vicinity of the active node. There are two generalizations of the 2D scheme: first, the remeshing process around an active node can be done for curved instead of flat surfaces, assuming a line-shaped scalpel. Second, a line-shaped scalpel can intersect a curved 3D surface in multiple points, so a consistent model of cutting allows multiple incisions, each with an active node. These active nodes can interact: incisions may meet, leading to annihilations, or incisions may hit folds, leading to branches.




Figure 5.7: Exits are predicted: when the extrapolated scalpel movement hits a boundary closely (a), then a new node is inserted at that location (b). The new edge is dissected (c) and contracted (d).

We assume that the surface is given as a triangle mesh with a boundary, and that no further information on the surface shape is known. The virtual scalpel is a line segment, and its movement is given by sampled positions of its endpoints. The endpoints are assumed to move with constant velocity between the samples. During a cut, active nodes are part of the boundary of the surface. The two boundary edges incident with an active node again are called cut edges, and define an incision.

5.3.1 Single incision

A scalpel movement is handled as follows: the line segment representing the new scalpel position is intersected with all active triangles. If a single intersection is found, then the active node is moved to that point. The active triangles are subjected to flipping. The 2D criterion is used to determine whether an edge is flipped: an edge is considered illegal when the two triangles incident to that edge would be illegal in a 2D triangulation. Conceptually, we could say that the incident triangles are unfolded to be coplanar, and then the two-dimensional criterion is used. It is not clear whether this flipping criterion leads to a terminating algorithm when applied to the edges of an arbitrary 3D surface mesh. In this sense, this technique is now truly a heuristic.

Flipping on 3D surfaces is a delicate operation: some flips are topologically impossible (as demonstrated in Figure 5.9). This means that the all operations must be checked for failure cases.

Incisions are split analogously to the 2D case: during a cut, the last path of the scalpel is stored. When the cut edges violate the 3D empty-circle criterion, new nodes are inserted where the path is closest to the line connecting the nodes opposite the cut edges.




Figure 5.8: Delaunay cutting on a triangle mesh. Both the mesh itself and the boundary are shown. The scalpel trajectory is indicated with a dotted curve. Notice how strongly curved path segments are cut short in the realized cut. The starting mesh was regular and consisted of 722 triangles. The cut increased mesh size by only 57 triangles.




Figure 5.9: Not all flips in 3D are topologically valid: when flipping the bold edge on the left, the new edge pair (bold) occurs twice in the resulting complex. The triangles are displayed in an exploded view for clarity.

Node removal is done in a heuristic manner. Suppose that we want to remove a node v that is incident with n triangles. We say that two adjacent triangles incident with v form an *ear*, if the sum of the angles opposite v is less than π . When removing v , all ears are flipped until v is incident with only 3 triangles. Then v is removed, and the involved edges are flipped back if they violate the empty-circle criterion. Ears always exist: suppose that the triangles incident with v are numbered $i = 1, \dots, n$, and have angles $\alpha_i, \beta_i, \gamma_i$, where γ_i is the angle at v (See Figure 5.10). Then there must be an $\alpha_{(i+1) \bmod n} + \beta_i < \pi$, since $n\pi = \sum_i (\alpha_i + \beta_i + \gamma_i) = \sum_i (\alpha_{(i+1) \bmod n} + \beta_i) + \sum_i \gamma_i$, and $\sum_i \gamma_i > 0$.




Figure 5.10: When removing a node, ears are flipped until v is in only 3 triangles. An ear has $\alpha_{(i+1) \bmod n} + \beta_i < \pi$, so the triangles incident with vw form an ear.

When the scalpel is almost parallel to the surface, small movements of the scalpel can result in large movements of an active node. For this reason it is necessary to control large movements. Large movements are subdivided using a maximum distance. This distance is computed as the minimum for all distances between the line spanned by the scalpel and lines spanned by the edge opposite the active node, as indicated in Figure 5.11. If the end-points of the scalpel segment cannot move further than d_{\max} as indicated in Figure 5.11, then the scalpel will not hit that edge. This can be shown as follows: suppose that the scalpel position is given by the line segment ht for some $h, t \in \mathbb{R}^3$, and the points corresponding to h and t move by amounts p and $q \in \mathbb{R}^3$ respectively, with $\|p\|, \|q\| < d_{\max}$. If the movement puts the scalpel on the line spanned by the edge, then $\lambda(p + h) + (1 - \lambda)(q + t)$ is on the edge indicated for some $0 \leq \lambda \leq 1$. By the definition of d_{\max} we have

$$d_{\max} \leq \|(\lambda(p + h) + (1 - \lambda)(q + t)) - (\lambda h + (1 - \lambda)t)\|$$

On the other hand, the latter expression equals $\|(\lambda\mathbf{p} + (1 - \lambda)\mathbf{q})\|$, which is bounded by $\lambda\|\mathbf{p}\| + (1 - \lambda)\|\mathbf{q}\| < d_{\max}$. This is a contradiction, so the movement (\mathbf{p}, \mathbf{q}) can not hit the edge.




Figure 5.11: The maximum movement d_{\max} for a scalpel and a single triangle. The total maximum distance is given by the minimum over all triangles incident with the active node.

5.3.2 Multiple incisions

In 3D, the scalpel can interact with the entire surface, and the scalpel may enter the mesh in any place. We can distinguish three cases, as demonstrated in Figure 5.12: the scalpel hits a boundary edge of the mesh, the scalpel tip enters through the surface, or the scalpel hits an internal edge of the mesh. The first case is handled completely analogous to the 2D case. The rest of the cases do not occur in 2D.




Figure 5.12: There are three ways for the scalpel to enter: by hitting the boundary, entering with the tip or hitting an exposed internal edge.

In the second case, a new active node is inserted in the incised triangle. When the active node is far enough from the entry point, a node is inserted at the original entry point. The result is an edge that connects the entry node with the active node. This edge is changed into a real incision, when the split action (shown in Figure 5.5) is executed. Until that time, the edge is constrained, so it cannot be flipped away. The third case is when the scalpel hits an exposed edge of the mesh. Then a single active node is inserted in the edge. During the next movement, the cut will branch into two incisions.

The scalpel is represented by a line segment, and line segments can interact with curved surfaces in many places: the scalpel may incise the surface in multiple locations, and during a cut a single incision may branch into multiple incisions, as is shown in Figure 5.14. We could apply the 3D remeshing process from Subsection 5.3.1 if we could rule out any interactions between different incisions. Fortunately, this seems possible: we can forbid interactions by ensuring that every triangle is incident with at

most one active node. This is achieved by the following restrictions on the meshing process.

- Incisions in edges or triangles that already active are rejected.
- Nodes that separate active nodes cannot be removed.
- Edges that separate two active nodes cannot be flipped.

When an active node moves towards a restricted node, an annihilation is performed: all edges from the restricted node to an active node are dissected. This is demonstrated in Figure 5.13. When an active node comes close to an edge separating active nodes, then its movement is extrapolated, and a new node is inserted at the extrapolated point. The annihilation now proceeds with the newly inserted node. A normal boundary exit is a special case of an annihilation.




Figure 5.13: Edges that separate incisions may not be flipped, nodes separating incisions may not be removed. Instead, when an active node comes too close to such a forbidden node or edge, an annihilation is performed (right).




Figure 5.14: Folded surfaces may lead to branching cuts (top). In such cases, a movement will cause the scalpel to intersect multiple triangles (bottom; intersection points are indicated by dots). When this happens, multiple incisions replace the old active node.

When a scalpel movement is processed, the next position of the active node is determined by intersecting all triangles incident with the new scalpel position. If multiple active triangles are intersected, as demonstrated in Figure 5.14, then the cut will branch. New incisions are inserted, the old incision is marked as no longer active, and the edges

connecting old and new active nodes are dissected. During a branch, the new nodes are close to the original incision. This leads to short edges and degenerate triangles. When the scalpel progresses further, these short edges disappear.




Figure 5.15: Intersecting a line sweep with an edge.

Collisions are computed as scalpel/edge intersections. The procedure for computing such intersections is as follows. We assume that the scalpel moves from ht to $h't'$ for some $h, t, h', t' \in \mathbb{R}^3$, and that we want to compute intersections with the line segment ab , for $a, b \in \mathbb{R}^3$ (See Figure 5.15). In other words, we want to solve

$$\begin{aligned} \mathbf{p} &= \gamma \mathbf{h} + (1 - \gamma) \mathbf{h}' \\ \mathbf{q} &= \gamma \mathbf{t} + (1 - \gamma) \mathbf{t}' \\ \lambda \mathbf{p} + (1 - \lambda) \mathbf{q} &= \mu \mathbf{a} + (1 - \mu) \mathbf{b} \\ \lambda, \gamma, \mu &\in [0, 1] \end{aligned}$$

Substituting the first two equations in the last one yields the equation

$$\lambda \gamma ((\mathbf{h} - \mathbf{t}) - (\mathbf{h}' - \mathbf{t}')) + \lambda (\mathbf{h}' - \mathbf{t}') + \mu (\mathbf{b} - \mathbf{a}) + \gamma (\mathbf{t} - \mathbf{t}') = \mathbf{b} - \mathbf{t}'.$$

This may be seen as a linear system in four variables. If we set $x = (\lambda \gamma, \lambda, \gamma, \mu)$, $\mathbf{c} = (\mathbf{h} - \mathbf{t}) - (\mathbf{h}' - \mathbf{t}')$ and write the 3×3 matrix \mathbf{A} for $(\mathbf{h}' - \mathbf{t}', \mathbf{t} - \mathbf{t}', \mathbf{b} - \mathbf{a})$, then we can rewrite the equation as

$$(\mathbf{c}|\mathbf{A})x = \mathbf{b} - \mathbf{t}', \quad (5.1)$$

where $x \in \mathbb{R}^4$ is the unknown. When viewed as a linear system, a solution may be given as $\mathbf{y} + \alpha \mathbf{k}$, where $\mathbf{y} \in \mathbb{R}^4$ is a particular solution of the system, and $\mathbf{k} \in \mathbb{R}^4 \setminus \{0\}$ is in the kernel of $(\mathbf{c}|\mathbf{A})$. We set $\mathbf{k} = (1, -\mathbf{A}^{-1}\mathbf{c})$ and $\mathbf{y} = (0, \mathbf{A}^{-1}(\mathbf{b} - \mathbf{t}'))$. Using these values, we may derive a quadratic equation for α from Equation (5.1). We thus obtain triples (λ, γ, μ) that define intersections.

5.4 Results

The 2D version of this algorithm has been implemented in a small application written in Python. The mouse controls a virtual scalpel that can perform cuts in uniform meshes on a square grid. A sample is shown in Figures 5.8 and 5.16.

Fix-up, node removal and exits depend on points being close to or far from the active node. These notions are expressed in global constant thresholds for the distance. These thresholds are denoted by ϵ_{edge} for node removal, ϵ_{entry} for entry fixup, and ϵ_{exit} for predicting exits. They were set to a fraction of the average global edge length \bar{h} . For Figure 5.8 we used $\epsilon_{\text{edge}} = \bar{h}/3$, and for Figure 5.16 we used $\bar{h}/3$ and $\bar{h}/8$. The entry and exit tolerances were set to $\bar{h}/6$. This choice is somewhat arbitrary, but if ϵ_{edge} is set larger than $\bar{h}/2$, mesh complexity will be *reduced* during a cut. Lower values set a threshold for edge length. This does not directly control the accuracy of the cut trajectory represented in the mesh: the example in Figure 5.16 shows that lowering the threshold can decrease accuracy.

In Figure 5.16 the new cutting scheme for 2D is compared with a subdivision approach, where each sliced triangle is replaced by three triangles. Element shapes are much worse in the subdivision mesh, as evidenced by the histograms of minimum and maximum triangle angle in Figure 5.17. On the other hand, the subdivision cut follows the scalpel path more closely.

The 3D version of this algorithm has been implemented in a small application written in C++. It also uses the mouse to obtain scalpel movements. Samples are shown in Figures 5.18 and 5.19. The 3D version assumes that all events are strictly ordered in time, and does not take special precautions for degenerate cases. When the scalpel enters or exits in a movement parallel to the surface, multiple events happen simultaneously, which leads to various failures.

Flips can be implemented in constant time and node removals can be accomplished in $\mathcal{O}(d \log d)$, where d is the degree of the node [34]. The expected value of d in surface meshes is 6, so the total remeshing process for a single movement can be implemented in constant time, which guarantees that the remeshing process itself is scalable to larger meshes. However, in the 3D version the scalpel can interact with all parts of the mesh at every step. In practice, efficient collision detection will be needed for a scalable implementation. Also, the number of events caused by a single scalpel movement depends on the angle between the surface and the scalpel. However, for small meshes (say, 1000 triangles) response of the 3D version is instantaneous when scalpel is more or less perpendicular to the surface (on a P3/1Ghz).

5.5 Discussion

We have presented an approach to cutting in triangulations that produces measurably smaller and better-shaped meshes than subdivision methods. The method is based on a model of a point-shaped scalpel that moves an active node through a static mesh. During movements, the area around the active node is remeshed. The approach generalizes curved surfaces in 3D, where the scalpel is line-shaped. The technique bears some resemblance to interactive mesh dragging, a technique proposed by Suzuki et al. [90] to




Figure 5.16: The same cut performed with both subdivision (top) and Delaunay cutting (center, $\varepsilon_{\text{edge}} = \bar{h}/3$ and bottom, $\varepsilon_{\text{edge}} = \bar{h}/8$). The starting mesh contains 50 triangles. The subdivision cut increases size by 62 triangles, the Delaunay cut in the center by 10, and at the bottom by 12 triangles. The scalpel trajectory is indicated with a dotted line. Notice that the mesh in the center picture matches the trajectory better than the mesh at the bottom, although $\varepsilon_{\text{edge}}$ was smaller.




Figure 5.17: Histogram of minimum (top) and maximum (bottom) element angle for the subdivision and $\bar{h}/3$ Delaunay cut from Figure 5.16. The peaks correspond to the unaltered triangles (minimum angle $\pi/4$, maximum angle $\pi/2$). The Delaunay cut (right bars) yield less extreme element angles than the subdivision cut.

allow dragging operations on triangle surfaces in interactive geometric modelers. It is possible to generalize the single-incision 3D technique to multiple incisions consistently.

The technique that we have described is heuristic: it employs more or less arbitrary constant tolerances ϵ_{edge} , ϵ_{exit} and ϵ_{entry} . This implies that it is only usable for meshes with a uniform resolution, and that suitable values must be computed beforehand. In retrospect, it might have been wiser to consider an edge too short when the opposite angle in incident triangles is shorter than some threshold. This criterion is also local but scale independent.

We have used a static model for mesh cutting. In the context of surgery simulation and deformable models, this is not realistic. Surgical instruments always interact with the deformed mesh, and exert force on the tissue that is cut. A full-fledged simulation would be able to respond to movements with reaction forces, which could be relayed to a force-feedback device. Deformation could be handled as follows: both the 2D and 3D algorithm include a step where the new position of the scalpel is intersected with parts of




Figure 5.18: In principle, our approach also works for a cut branching into three cuts. The mesh has very sharp angles and a limited resolution, so edge flips produce an odd end result.

the mesh. When deformation is present, these intersections are done with the deformed mesh. The deformation of the triangles involved can be used to translate the point back to the reference situation. Our technique can then proceed as described using reference locations. However, this does not address a more fundamental limitation: our model assumes that a more or less uniform mesh is desirable. It can be expected that more accuracy for the deformations are required close to the scalpel, since this is where forces are exerted on the material. When accurate deformations are needed, the mesh should be adaptively refined close to the scalpel incision, and coarsened further away.

Entry and exit are also changed by the presence of deformation: surgical instruments must overcome a threshold in force to puncture membranes covering organs [15, 36]. Cutting requires less force than puncture, so after the instrument enters the tissue, it will make an incision immediately, ensuring that entry does not lead to arbitrarily small incision depths. This is a physical variant of our entry-fixup, and it would make our own entry-control superfluous. A similar observation might also hold for a scalpel finishing a cut. Removing entry control is attractive, since our approach depends on a heuristically chosen quantity ϵ_{entry} , and it is possible to create a deadlock situation: if a scalpel entry has to be fixed up, the active node is part of the boundary. It is not always possible to preclude element inversion when moving boundary nodes, so the step-size control for large movements may halt further movements in some cases.

During a cut, nodes are added and removed; the nodes are added on line segments connecting existing nodes, which implies that the overall resolution of the mesh does not increase. Some interactive simulations of deformable objects refine meshes on demand to provide more accurate results in the region of interest [36, 76, 100]. Delaunay refinement algorithms [25, 82] seem to fit our framework of using Delaunay Triangulations, however more research is needed before this can be used in practice. Refinement algorithms need input geometries without small angles. Moreover, in 3D a surface triangulation is an approximation of a smooth surface. It is necessary to know the original




Figure 5.19: Screenshots from the 3D prototype, showing the cuts effected with mouse movements. The mesh is a 392 triangle surface mesh of a Gaussian bell-curve. The first 5 images show the evolution of two cut movements. The final image demonstrates more self-intersecting paths.

surface for determining where to insert new vertices; existing 3D surface meshing algorithms, such as Chew's guaranteed quality surface Delaunay refinement [25], also assume that a such smooth description is available. Hence, a cutting algorithm using refinement should be redesigned with the assumption that a surface shape itself is known.

We have demonstrated an extension of our single incision 3D approach to multiple incisions. The extension is consistent with our model of the scalpel as a moving line segment. It remains to be seen how branching cuts should be combined with deformations.

Unfortunately, both the rationale for using Delaunay triangulation and our heuristics do not readily generalize to tetrahedral meshes. Delaunay tetrahedralizations of well-spaced points admit *slivers*, tetrahedrons that have four nearly planar and nearly cocircular vertices (see Figure 5.20). Slivers are flat tetrahedrons and therefore undesirable. Moreover, the higher-dimensional equivalent of edge flipping (face flipping), does not always work: there exist configurations of non-Delaunay faces which can not be flipped. This implies that flipping as a local improvement strategy does not always work. If the tetrahedralization is already Delaunay, then flipping non-Delaunay faces during insertion or removal of points in a mesh is always possible. Since cuts are non-convex, this suggests that a tetrahedral generalization must build a constrained or conforming Delaunay tetrahedralization with a moving boundary. A first step towards enabling such cuts would be to extend the current 3D surface cutting approach to cuts of closed surfaces: when cutting a closed surface, surface triangles are added to the inside of the incision, so the surface remains closed during cuts. This surface could then be used as a basis for making a tetrahedralization. An additional complication is that a constrained Delaunay tetrahedralization only exists if edges have no non-incident nodes close by (the surface should be *ridge-protected*) [86].




Figure 5.20: Slivers have an excellent circumcircle to shortest-edge ratio, so they can be present in Delaunay tetrahedralizations of well-spaced point sets. They are degenerate nevertheless.