Pacific Journal of Mathematics

SIMPLIFYING TRIANGULATIONS OF S^3

Aleksandar Mijatović

Volume 208 No. 2 February 2003

PACIFIC JOURNAL OF MATHEMATICS
Vol. 208, No. 2, 2003

SIMPLIFYING TRIANGULATIONS OF S^3

Aleksandar Mijatović

In this paper we describe a procedure to simplify any given triangulation of S^3 using Pachner moves. We obtain an explicit exponential-type bound on the number of Pachner moves needed for this process. This leads to a new recognition algorithm for the 3-sphere.

1. Introduction.

It has been known for some time that any triangulation of a closed PL n-manifold can be transformed into any other triangulation of the same manifold by a finite sequence of moves [**5**]. We can describe the moves as follows.

Definition. Let T be a triangulation of an n-manifold M. Suppose D is a combinatorial n-disc which is a subcomplex both of T and of the boundary of a standard $(n+1)$-simplex Δ^{n+1}. A *Pachner move* consists of changing T by removing the subcomplex D and inserting $\partial\Delta^{n+1} - \mathrm{int}(D)$ (for n equals 3, see Figure 1).

It is an immediate consequence of the definition that there are precisely $(n+1)$ possible Pachner moves in dimension n. We can now state Pachner's result [**5**] in the following way.

Theorem 1.1 (Pachner)**.** *Closed* PL *n-manifolds M and N with triangulations T and K respectively, are piecewise linearly homeomorphic if and only if there exists a finite sequence of Pachner moves and simplicial isomorphisms taking the triangulation T into the triangulation K.*

In dimension 3 we have four moves from Figure 1 at our disposal. Using them, we can describe the main theorem of this paper.

Theorem 1.2. *Let T be a triangulation of a 3-sphere and let t be the number of tetrahedra in it. Then we can simplify the triangulation T to the canonical triangulation of S^3, by making less than $a\, t^2 2^{bt^2}$ Pachner moves, where the constant a is bounded above by $6 \cdot 10^6$ and the constant b is smaller than $2 \cdot 10^4$.*

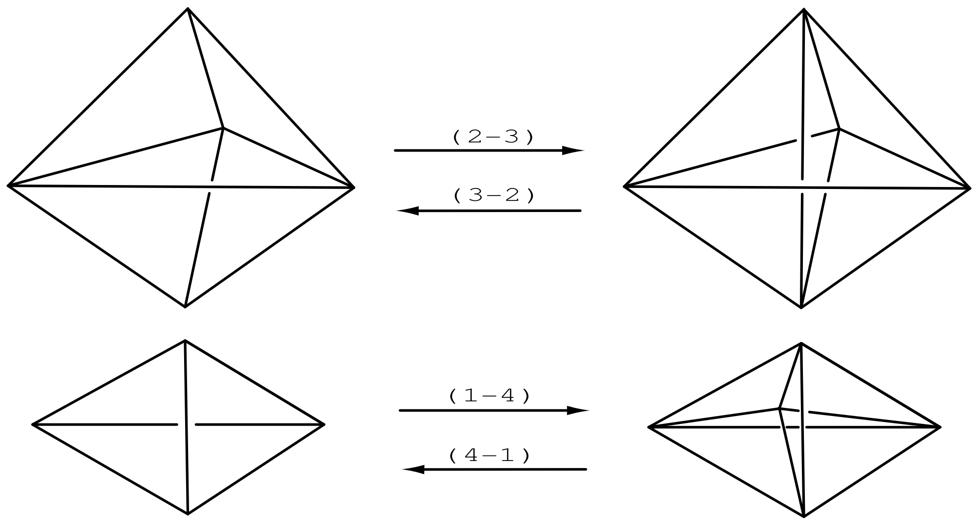

Figure 1. Three dimensional Pachner moves.

The triangulation T in this theorem can be non-combinatorial (i.e., simplices are not uniquely determined by their vertices), as is the case with the canonical triangulation of S^3, consisting of two standard 3-simplices glued together via an identity on their boundaries. We should mention here that Pachner's original proof of Theorem 1.1 works for combinatorial triangulations only. However, at least in dimension 3, this does not matter because the second derived subdivision of any (possibly non-combinatorial) triangulation is always combinatorial and can be obtained from the original triangulation by a finite sequence of Pachner moves.

A possible effect Pachner's result could have on the theory of 3-manifolds is discussed by the next proposition.

Proposition 1.3. *Let T and K be two triangulations of the same closed* PL 3*-manifold M. The existence of a computable function, depending only on the number of* 3*-simplices in T and K, bounding the number of Pachner moves required to transform T into K, is equivalent to an algorithmic solution of the recognition problem for M among all* 3*-manifolds.*

Proof. Assume first that $f(t, k)$ is a computable function as described in the proposition. Suppose that T is a triangulation of M with t 3-simplices. Let K be a triangulation of some closed 3-manifold N containing k 3-simplices. Do all possible sequences of Pachner moves on the triangulation T of length at most $f(t, k)$, and check each time if the result is isomorphic to K. This gives an algorithm to determine whether M and N are PL homeomorphic.

Conversely suppose that we have an algorithm to recognize M among all 3-manifolds. Now we need a complete (finite) list of all triangulations of all 3-manifolds with a fixed number of 3-simplices. In dimension three, such a list can be built algorithmically because there is an easy way of recognizing the 2-sphere (the Euler characteristic suffices) as a link of a vertex.

We can now create all triangulations of M with the specific number of 3-simplices by running the recognition algorithm for M (which exists by assumption) on the list of all 3-manifold triangulations with the specified number of 3-simplices.

An algorithm, making all possible Pachner moves on a triangulation of our 3-manifold M with t 3-simplices will after a finite number of steps (by Theorem 1.1) necessarily produce a given triangulation of M containing k 3-simplices. Since we can list all triangulations of M with t (respectively k) 3-simplices, this gives an algorithm to calculate the value of the function $f(t,k)$ as required. □

At present there is no known algorithm to decide whether a given simplicial complex is an n-sphere, for $n \geq 4$. This means that the proof of one of the implications in Proposition 1.3 breaks down in dimensions five and above since there is no way of building a list of all triangulations of all manifolds with a fixed number of top dimensional simplices, in these dimensions.

The proof of the converse implication in Proposition 1.3, showing that a computable bound implies a recognition algorithm for a given n-manifold, remains valid in any dimension. Furthermore, if such a computable bound existed for all n-manifolds, and was independent of the underlying n-manifold, then it would give an algorithm to determine whether any two n-manifolds are homeomorphic. But using the fact (proved by A.A. Markov) that there is no such algorithm for $n \geq 4$, we can conclude that such a computable function does not exist in dimensions four and above.

It is interesting to note, that for any n-manifold M Pachner's theorem implies the existence of a function, depending only on the number of n-simplices in T and K, and bounding the number of Pachner moves necessary for the whole transformation. This is because there are only finitely many triangulations of our n-manifold M with fixed numbers of n-simplices. Then, using Theorem 1.1, a finite sequence of Pachner moves connecting any two of them, can be found. Taking the maximum length over this finite family of sequences gives us the bound. Therefore, computability of the function in Proposition 1.3 is an assumption that can not be omitted.

The upper bound in Theorem 1.2 is computable. It therefore yields a new recognition algorithm for the 3-sphere.

2. Normal surfaces.

One of the essential ingredients of the proof of Theorem 1.2 is the theory of normal and almost normal surfaces. In this section we shall describe some of its basic features. We will then go on to discuss the Rubinstein-Thompson algorithm [**7**] for recognizing the 3-sphere which provides the setting for the proof of Theorem 1.2. After it, we'll mention some of the consequences of normal surface theory which will prove to be useful later. At the end of this section we shall prove the isotoping lemma that will later give us a way of simplifying triangulations of the 3-sphere. Let's start with some definitions.

A *normal triangle* (respectively *quadrilateral*) in a 3-simplex Δ^3 is a properly embedded disc D, such that its boundary ∂D intersects precisely three (respectively four) edges transversely in a single point and is disjoint from the remaining 1-simplices and vertices of Δ^3. A *normal disc* is a normal triangle or quadrilateral.

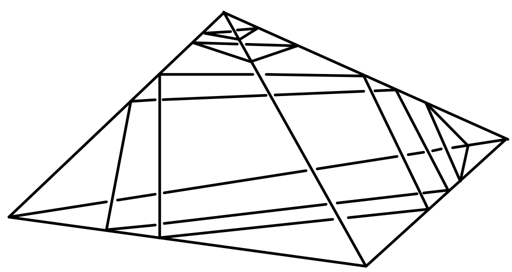

Figure 2. Three types of normal discs.

There are four possible types of normal triangles, because each triangle is parallel to one of the faces of Δ^3. Normal quadrilaterals will always separate the vertices of the tetrahedron in pairs. It is therefore clear, that we can only have three possible quadrilateral types. Together, there are 7 distinct normal disc types in a tetrahedron.

Let M be a 3-manifold with a triangulation T. A properly embedded surface F in M is in *normal form* with respect to the triangulation T, if it intersects each tetrahedron of T in a finite (possibly empty) collection of disjoint normal discs.

Since normal surfaces are always embedded, at most one of the quadrilateral types can occur in each 3-simplex.

Suppose F is a normal surface in M with respect to T. Then F corresponds to a vector $\underline{x} = (x_1, \ldots, x_{7t})$ with $7t$ coordinates, where t denotes the number of 3-simplices in the triangulation T. The index set $\{1, \ldots, 7t\}$ corresponds to all possible disc types in T (there is 7 of them for each tetrahedron). The coordinate x_i is simply the number of copies of i-th disc type in our surface F.

Each 2-simplex in T contains three types of normal arcs (coming from normal discs), one cutting off each vertex of the triangle. If it is a face of two 3-simplices in T, then it gives rise to three matching (linear) equations, one corresponding to each normal arc type. Doing this for every triangle, not in the boundary of M, we've constructed a linear system in $7t$ variables, consisting of at most $6t$ equations.

It follows immediately from the construction, that the vector $\underline{x}$, coming from the normal surface F, gives a solution to the linear system. By imposing extra conditions to ensure that all quadrilaterals in a given tetrahedron are of the same type, we obtain a restricted linear system. The conditions we've just added are sometimes referred to as *quadrilateral constraints.* Now there is a one to one correspondence between embedded normal surfaces in M and nonnegative integral solutions to the restricted linear system.

Haken proved that all nonnegative integral solutions to such a system are integer linear combinations of a finite set of nonnegative integral solutions $\underline{x_1}, \ldots, \underline{x_n}$, called *fundamental solutions*, which can be found in an algorithmic way. As it turns out, these fundamental solutions are characterized by the property of not having a decomposition as a sum of two (nontrivial) nonnegative integral solutions to the restricted linear system.

Since each fundamental solution corresponds to an embedded normal surface, we obtain a finite set $F_1, \ldots, F_n$ of embedded normal surfaces, called *fundamental surfaces.* Any embedded normal surface in M can thus be written algebraically as a nonnegative integer linear combination of fundamental surfaces. Miraculously, this algebraic fact carries over to geometry. In other words, we can define a geometric addition for any two normal surfaces F and G with the property that the sum of the corresponding solutions to the restricted linear system, is again a solution of the same system. This condition boils down to the fact that the union of all normal discs in both F and G satisfies the quadrilateral constraints.

Assuming that and putting both surfaces in general position with respect to one another, cutting along the arcs of intersection in each tetrahedron, and pasting the pieces back together in the unique way, so that we end up with normal discs only, yields a well-defined embedded normal surface $F+G$. Its corresponding vector is a sum of the vectors coming from F and G. The cut and paste process described above is sometimes called *regular alteration.*

An isotopy of the ambient manifold, preserving the normal structure of a given normal surface is called a *normal isotopy.* We should also note that the geometric addition described above is well-defined up to a normal isotopy of the summands.

Before we describe the Rubinstein-Thompson algorithm, we need to introduce a concept, originally due to Rubinstein.

Definition. A properly embedded surface in a 3-manifold M with a triangulation T is *almost normal* with respect to T, if it intersects each tetrahedron of T in a finite (possibly empty) collection of disjoint normal discs except in precisely one tetrahedron there is precisely one exceptional piece from Figure 3 and possibly some normal triangles.

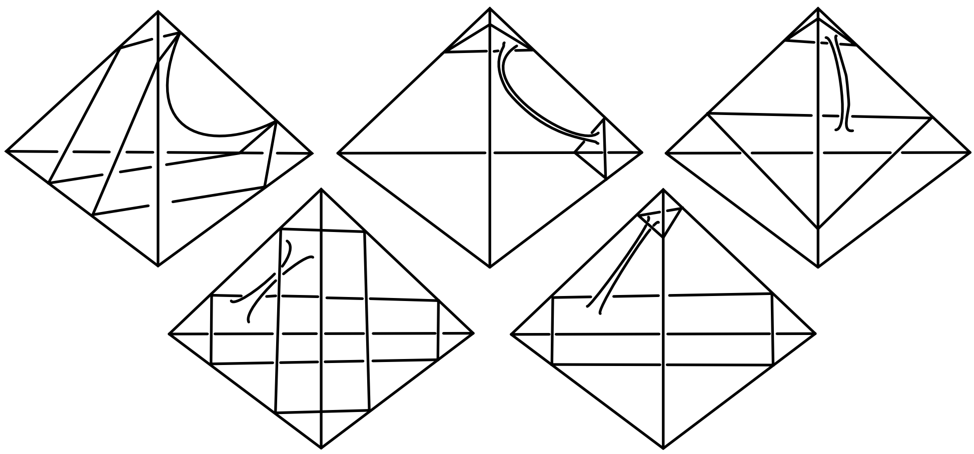

Figure 3. Almost normal pieces.

This exceptional piece is either a disc (the first possibility in Figure 3) whose boundary is a normal curve of length eight (i.e., an octagon), or it is an annulus consisting of two normal disc types with a tube between them that is parallel to an edge of the 1-skeleton.

Now we can describe the Rubinstein-Thompson algorithm which is designed to determine whether or not a 3-manifold M with a triangulation T is a 3-sphere. We can assume that M is closed, orientable and that $H_1(M;\mathbb{Z}_2)$ is trivial. All these properties can be checked algorithmically. The last assumption guarantees that M contains no closed non-separating surfaces. The algorithm now is in three steps. We proceed as follows.

Step 1. Find a maximal collection Σ of disjoint non-parallel normal 2-spheres in M.

Step 2. Cut M open along Σ. This splits M into three different types of pieces:

Type A: A 3-ball neighborhood of a vertex of T (every vertex is enclosed in such a piece).

Type B: A piece with more than one boundary component.

Type C: A piece with exactly one boundary component which is not of Type A.

Step 3. Search each Type C piece for an almost normal 2-sphere with an octagonal component.

Conclusion: M is a 3-sphere if and only if every Type C piece contains an almost normal 2-sphere with an octagonal component.

The bulk of the proof that this indeed is a recognition algorithm for the 3-sphere relies on the following two lemmas from [**7**].

Lemma 2.1. *A Type* B *piece is a punctured* 3*-ball.*

Lemma 2.2. *A Type* C *piece is a* 3*-ball if and only if it contains an "octagonal" almost normal* 2*-sphere.*

By Lemma 2.2, if some Type C piece fails to contain an "octagonal" almost normal 2-sphere, then it is not a 3-ball and M is not a 3-sphere. Otherwise, M is just a collection of 3-balls and punctured 3-balls glued together. Since every 2-sphere is separating, M has to be a 3-sphere.

The difficult part of the argument is in the proof of Lemma 2.2. It is here that Thompson simplified Rubinstein's original methods to prove the existence of an "octagonal" almost normal 2-sphere in a 3-ball of Type C, by using Gabai's powerful notion of thin position. We should also note that the easier converse implication in Lemma 2.2 follows from Lemma 2.7.

In order to be able find a maximal collection of disjoint non-parallel normal 2-spheres in M in an algorithmic way, we need the following lemma.

Lemma 2.3. *A maximal collection* Σ *of disjoint non-parallel normal* 2*-spheres in* M*, as in the Rubinstein-Thompson algorithm, can always be constructed algorithmically.*

A proof of this lemma was given by Casson [**1**]. It was also described in [**3**] (see Lemma 3). It will be important for us to be able to bound the complexity of all of the 2-spheres in the maximal collection Σ. We shall therefore give a brief description of this algorithm.

Additivity of Euler characteristic implies at once that if there exists a nontrivial normal 2-sphere in our triangulation, we can also find one (which is also nontrivial) among fundamental surfaces. Since the family of fundamental surfaces is accessible in an algorithmic way, we can take this fundamental 2-sphere to be the first element in Σ.

Assume now that we have already constructed a subcollection Σ' of Σ. We shall look for normal surfaces with respect to the triangulation T, lying in the complement of the normal 2-spheres constructed so far. In any tetrahedron from T we can have complementary regions of Σ' that are not of the form triangle$\times I$ or square$\times I$ (see Figure 4), as well as the ones that are. Note also that the unions of the product regions support a natural I-bundle structure. These I-bundles are usually referred to as parallelity regions.

We can now describe normal surfaces in the complement of Σ' by assigning a variable to each triangle or square type that does not lie in any of the parallelity regions. The equations are again the matching equations along the faces together with the equations that ensure that the surface intersects each parallelity region in a well-defined number of components. In other words each parallelity I-bundle contributes new linear equations (that are of the same form as the matching equations along faces) and no new variables. This is because the normal 2-spheres we are looking for, can only run parallel to the horizontal boundary in these I-bundles. Adding the usual quadrilateral constraints gives a restricted linear system.

Like in standard normal surface theory, we have a one to one correspondence between closed normal surfaces in the complement of Σ' and solutions of the above restricted linear system. Since addition of two solutions is again realized topologically by regular alteration, the same argument as before tells us that the next normal 2-sphere, that is not parallel to any of the elements in Σ', can be chosen from the family of fundamental surfaces.

So in order to find a maximal family Σ of disjoint non-parallel normal 2-spheres, we just have to keep repeating this procedure. We stop when each normal 2-sphere in the complement of Σ' is normally parallel to some normal 2-spheres in the collection Σ'. Lemma 2.4 guarantees that this process has to reach such a stage.

As far as the complexity, i.e., the number of normal pieces, of elements in Σ goes, at each stage it is going to be bounded by Proposition 2.5. Since the linear algebra in the proof of Proposition 2.5 (which can be found in [**2**]) depends only on the number of normal variables and is independent of how many equations we have, it is the number of different normal disc types outside the parallelity regions, that needs to be controlled. The proof of Lemma 2.4 shows that this number is bounded linearly by the number of tetrahedra in T. In fact there can be at most 11 different normal disc types in a single tetrahedron in T at any stage of the process. We shall calculate explicit upper bounds later on in this section.

We still need to answer the question of how to search for "octagonal" almost normal 2-spheres that are contained in Type C pieces. Modified versions of standard normal surface theory algorithms suffice for the search. So our goal is to construct an algorithmic procedure which will find an "octagonal" almost normal 2-sphere in each Type C piece. These 2-spheres will exists by Lemma 2.2 if the 3-manifold M we are looking at is a 3-sphere. We proceed as follows.

First we fix a tetrahedron H in the triangulation T of a 3-manifold M and then we fix a normal curve c of length eight on its boundary (there are three choices for c). Now an analogue to the normal surface theory, used to construct the collection Σ, can be set up. The matching conditions will look just like before. Quadrilateral constraints have to be modified however, because we want our solutions to consist of normal triangles and quadrilaterals everywhere except in H, where we want them to be composed of normal triangles and octagonal components with boundaries parallel to c. The notion of regular alteration can be defined in this generalized setting and again it gives rise to the correspondence between integer linear combinations of the fundamental solutions to the (generalized) restricted linear system and the set of all surfaces described above. Fundamental surfaces are again the ones corresponding to fundamental solutions. We should also note that their complexity is bounded by Proposition 2.5 since the linear system they are the solutions of, has less than $11t$ variables.

What we really want is to find algorithmically "octagonal" almost normal 2-spheres that are contained in Type C pieces. We know that one such 2-sphere exists in each Type C piece by Lemma 2.2. This 2-sphere can be expressed as a sum of the fundamental surfaces. Precisely one of the summands has to contain a single octagonal piece and, since the Euler characteristic is additive, at least one of the fundamental surfaces in the sum has to be a 2-sphere (since the Type C piece we are looking at contains an "octagonal" almost normal 2-sphere, it has to be a 3-ball and can therefore not contain embedded projective planes). If the fundamental 2-sphere in the sum does not contain an octagon, then it has to be normal and thus parallel to the unique normal 2-sphere from Σ that is bounding the Type C piece we are looking at. This is a contradiction because we could then isotope it away from all the other summands by a normal isotopy. Since regular alteration is defined up to normal isotopy, this would then make the sum (i.e., a 2-sphere) disconnected. So we've found an "octagonal" almost normal 2-sphere in a Type C piece that is fundamental.

The complexity of the fundamental "octagonal" almost normal 2-sphere we've just constructed is bounded in the same way as all the other complexities of the normal 2-spheres in Σ. This follows directly from the construction, since all we are doing when searching for an almost normal 2-sphere, is just making another step of the recursion that gave us Σ, without increasing the number of normal variables. We will give an explicit estimate for the complexity later on in this section.

Let's first bound the number of disjoint non-parallel normal 2-spheres in Σ. This is made possible by an old idea due to Kneser.

Lemma 2.4. *Let T be any triangulation of S^3 and let t be the number of tetrahedra in T. Then any family of disjoint non-parallel normal 2-spheres contains at most $6t$ of them.*

Proof. Normal triangles and squares chop up any tetrahedron in T into several pieces. But at most six of these regions are not of the form triangle$\times I$ or square $\times I$ (see Figure 4).

Let n be the maximal number of disjoint non-parallel normal 2-spheres in T. Then the complement of this family has precisely $(n+1)$ components. Each of those components must contain at least one of the non-product

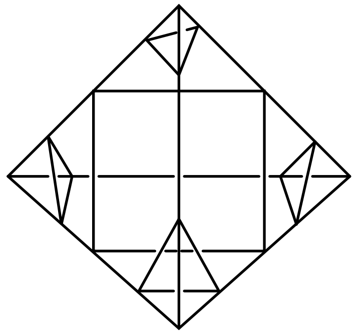

Figure 4. Complementary regions which are not a product.

regions. This is because any component, consisting only of product pieces, is bounded by two parallel normal 2-spheres. Since the total number of non-product regions is bounded by $6t$, our lemma is proved. $\square$

We are interested in bounding the number of normal pieces of elements in Σ. We also want to bound the number of normal pieces of the "octagonal" almost normal 2-spheres that arise in Type C pieces. Both of these things can be accomplished at one go, because we know that the procedure giving Σ can be extended (by making a single additional step) to an algorithm producing "octagonal" almost normal 2-spheres in Type C pieces.

The proposition we are about to state is proved in [**2**]. It originally deals with the linear system in $7t$ variables coming from the matching equations for normal surfaces. Its proof uses some basic linear algebra on the linear system which consists of matching equations. We should note at this point that the number of equations in this linear system does not influence the bound that the proposition gives.

Proposition 2.5. *Let M be a triangulated 3-manifold containing t tetrahedra. Let $\underline{x}$ be a fundamental solution of a system of linear equations coming from matching conditions. Then each coordinate of the vector $\underline{x}$ is bounded above by $7t2^{7t-1}$.*

Using Proposition 2.5, we can bound the size of each component of all the vectors corresponding to the normal 2-spheres in Σ. It follows from Figure 4 that the number of normal discs that are not contained in the parallelity regions (at any stage of the algorithm producing the family Σ) is always bounded above by $11t$. The system of equations we are solving at each stage consists of the matching equations along the faces together with the equations that ensure that the surface intersects each parallelity region in a well-defined number of components. We should note here that the latter equations are of the same form as the matching equations.

Proposition 2.5 then implies that there can be at most $11t2^{11t-1}$ parallel copies of a given normal disc type in a complement of any subcollection Σ' from any stage of the algorithm. This can be deduced because the proof of Proposition 2.5 depends only on the number of variables and the shape of the equations, i.e., the number and size of nonzero coefficients in each equation. Since the number of variables has increased and the shape of the equations hasn't changed, we get the above bound by substituting $7t$ with $11t$ into Proposition 2.5.

Lemma 2.4 tells us that we'll never have to make more than $6t$ steps when constructing Σ. This means that each of the normal disc types in the complement of any subcollection Σ' can only give rise to less than $(2 \cdot 11t \cdot 2^{11t-1})^{6t}$ normal discs of the same type in the initial triangulation T. This is because at each stage of the algorithm the number of parallel copies of a fixed normal disc type is given by $11t2^{11t-1}$. We have to include the factor of 2 because in each parallelity region there are two normal variables contributing to the number of parallel copies of a given normal disc type in the initial triangulation.

We can obtain a similar kind of bound for "octagonal" almost normal 2-spheres. We only have to change the exponent from $6t$ to $(6t+1)$. This is because all these "octagonal" almost normal 2-spheres are just one step away (in our algorithm) from the normal ones (bounding Type C pieces) and at each stage they are described by fewer variables. For example, in our original triangulation they require $7(t-1)+4$ variables. So the bound in Proposition 2.5 applies.

Putting everything together and using the fact that $5t(11t2^{11t})^{6t+1} < 2^{110t^2}$, we get the following lemma (the factor $5t$ comes in because there are at most 5 different normal and almost normal pieces in each tetrahedron of T).

Lemma 2.6. *Let T be a triangulation of the 3-sphere which contains t tetrahedra. Then the number of all normal pieces contained both in all elements of Σ and in all "octagonal" almost normal 2-spheres from all Type C pieces is bounded above by 2^{110t^2}.*

We should note at this point that this is the only part of the bound in Theorem 1.2 which contains a quadratic expression in its exponent. If one could find both the "octagonal" almost normal 2-spheres in Type C pieces and the maximal family Σ among the fundamental solutions of linear systems that are based on the triangulation T, the bound in Lemma 2.6 would have a linear function (similar to the one in Proposition 2.5) in its exponent.

The essential process we are just about to describe, is the one of isotoping almost normal surfaces. It is going to provide a foundation for the simplifying procedure needed for the proof of Theorem 1.2.

Let F be a separating almost normal surface in a 3-manifold M with a triangulation T. Its *weight*, $w(F)$, is defined to be the number of points in the intersection of F and the 1-skeleton T^1. If F contains an octagon, a *natural isotopy* is the one pushing the surface over an edge which meets the length eight normal curve bounding the octagon in two points. There are two possible natural isotopies, depending on the component of $M - F$ we are pushing into. In case of other non-normal pieces (see Figure 3), a natural isotopy pushes the tube part of the annulus so that it encompasses one of the edges it is parallel to. As a result in both cases, we get a surface with its weight equal to $w(F) - 2$.

Notice that if we look at our almost normal surface F in the complement of the 1-skeleton of T, there is an obvious compression disc D for it, enveloping the edge we are isotoping over. The natural isotopy can then be realized by isotoping over the 3-ball bounded by D and the disc in F, bounded by ∂D.

The natural isotopy is only the first step in the process of isotoping almost normal surfaces. Everything else will be accomplished by a sequence of elementary isotopies. We can define them as follows.

Let A be a 2-simplex in T, containing a non-normal arc (Figure 5) which comes from intersecting A with an isotope of F.

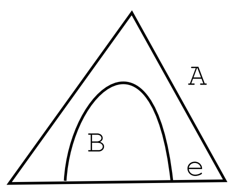

Figure 5. An isotope of F intersects A in a non-normal arc.

A disc B in the triangle A (see Figure 5) is bounded by the non-normal arc and a subarc of the edge e. An *elementary disc* can be constructed by banding together two parallel copies of B in the complement of 1-skeleton, where the band runs around the edge e. Its boundary is a simple closed curve in the surface, bounding a disc on one side. An *elementary isotopy* is an isotopy over the 3-ball bounded by the disc in the isotope of F and the elementary disc we've just defined.

Since F is a separating surface, we can fix a complementary component I of $M - F$. All the elementary isotopies that we are going to do from now on, are going to have the same direction. We will always be isotoping towards the interior of the component I.

The following isotoping lemma will play a crucial role in the simplifying process. A similar result is proved in [**6**] by a careful inspection of all the possible cases. The proof we are giving here is based on elementary isotopies and is better suited from our perspective because it sheds more light on the side of things we'll be interested in later.

Lemma 2.7. *Let F be a separating almost normal surface in a* 3*-manifold M with a triangulation T. Let I be a component of the complement $M - F$, if F contains an octagon. Otherwise let I denote the component containing a solid torus region in the interior of the* 3*-simplex where F is not normal. A natural isotopy followed by a sequence of all possible elementary isotopies, all going in the direction of I, will result with a surface intersecting each tetrahedron of T in pieces as in Figure* 6 *and in normal pieces. Moreover, in each tetrahedron there can only be at most one piece of the first type from Figure* 6*. A single* 3*-simplex can contain several pieces of all the other types in Figure* 6 *as well as several normal pieces. The pieces in Figure* 6 *can not be parallel.*

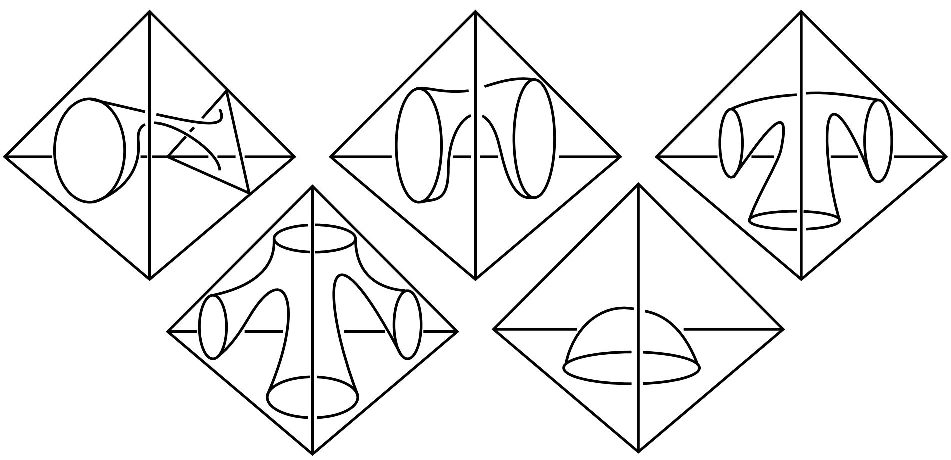

Figure 6. Non-normal pieces in the tetrahedra of T.

Proof. First note that after the natural isotopy, all the non-normal arcs we get will give rise to elementary isotopies in the direction of I. After each isotopy both F and I will change, but we'll still denote both resulting spaces by F and I respectively.

After the natural isotopy, F and I satisfy the following conditions:

1) In each tetrahedron of T the component I consists of a family of 3-balls, each one bounded by pieces of F and a (possibly disconnected) planar surface, contained in the boundary of the 3-simplex.
2) Each 3-ball from 1 intersects any face of the tetrahedron it lies in, in at most one disc.

An elementary isotopy moves a disc in F over a 3-ball in I, which intersects a single edge e in T^1. So the new I is just the old I without the 3-ball we isotoped over. This 3-ball is a union of a family of 3-balls, one in each tetrahedron of the star of e. In fact there can be more than one 3-ball from the same family in a single 3-simplex if this 3-simplex occurs more than once in the star of the edge e. This is perfectly feasible in a non-combinatorial triangulation, but it does not have any effect on the process we are studying.

The elements of the above family are the ones that are going to determine the topology of the pieces of I in tetrahedra of T. In fact, each 3-ball from condition 1 will after an elementary isotopy still satisfy both conditions if we substitute the old F with the new one. So after performing all possible elementary isotopies towards the interior of I, the surface F we end up with will intersect each triangle of T in normal arcs and simple closed curves which miss the boundary of the triangle.

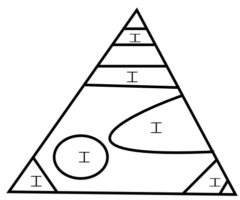

Figure 7. An intermediate state of the isotopy on a triangle in T.

The region I will, after the isotopy, consist of 3-balls in each 3-simplex. There is going to be a bijective correspondence between the 3-balls in the end, and the ones we started with. By condition 2, every 3-ball will still intersect any face of the 3-simplex it lies in, in at most one disc. It is also true that the number of these discs will not increase when we pass from the 3-ball pieces of I at the beginning to the 3-balls at the end.

Let's look at the pieces of F in each tetrahedron. It is obvious that all the possibilities of the lemma can actually arise. We have to see that they are the only ones.

Claim. A single piece of F can intersect a triangle of T in either a unique normal arc or in a single simple closed curve.

Every piece of F is contained in the boundary of a 3-ball piece of I. This 3-ball intersects each triangle of the 2-skeleton T^2 in at most one disc. So no triangle can contain two simple closed curves or a simple closed curve and a normal arc, both belonging to the same piece of F.

The same argument tells us that a triangle in T^2 can either contain two normal arcs of intersection with a single piece of F or at most three of them, each one cutting off a vertex of the triangle in the 2-skeleton.

Now we need to prove that our piece of F can have at most one normal simple closed curve boundary component. So assume the opposite. Since the piece is a subset of the boundary of a 3-ball, no arc contained in it, running between two distinct boundary components of our piece, can be extended to a simple closed curve in the 2-sphere bounding that 3-ball, without increasing the number of intersection points with the boundary of our piece. On the other hand, assuming we have at least two normal simple closed curves in the boundary, there surely exist two normal arcs, belonging to the distinct boundary components of our piece, that are contained in a single 2-simplex. Connecting them by an arc in the piece of F contradicts what was said before (because these two normal arcs are both contained in the boundary of a disc in the 2-simplex they lie in).

So now it follows that the piece of F we are looking at, can contain at most one normal boundary component which is of length at most eight. This is because the only normal curve of length 12, intersecting each 2-simplex (in the boundary of a tetrahedron) in 3 normal arcs as above, consists of 4 simple closed curves (one for each vertex of the tetrahedron). It is also well-known that normal simple closed curves of lengths 9, 10 or 11 do not exist.

There are precisely three normal simple closed curves of length eight in the boundary of a tetrahedron. So if our piece of F is bounded by one such curve, then at least one of the faces of the tetrahedron intersects the 3-ball piece of I (containing in its boundary the piece of F we are considering) in two discs. This is a contradiction that proves the claim.

The claim implies the following seven possible boundaries for any piece of F: Normal simple closed curve of length three, normal simple closed curve of length four, single simple closed curve, normal simple closed curve of length three and a simple closed curve, two simple closed curves, three simple closed curves, four simple closed curves.

Since every piece of F in any tetrahedron is planar, it is up to homeomorphism determined by its boundary. This implies that all possible pieces of F are the ones listed in the lemma.

The fact that all these planar surfaces are embedded as in Figure 6 (up to an isomorphism of the tetrahedron) follows from the observation that all the elementary discs are parallel to edges of the 1-skeleton. $\square$

3. Outline of the proof.

Given a triangulation T of the 3-sphere, how do we simplify it? The process is divided into two stages. First, we create a subdivision S of T by defining it in each complementary piece of the manifold $S^3 - \Sigma$ in such a way that the triangulations match along all normal 2-spheres in Σ. The second step consists of simplifying S down to the canonical triangulation of S^3.

An explicit construction of S, using Pachner moves, will be given in Section 5. The simplifying procedure of step two is based on the relationship between Pachner moves and shellable triangulations. This relationship will be established in Section 4.

Now, we are going to describe the additional structure on the complementary pieces of $S^3 - \Sigma$, needed for the definition of the subdivision S. We already know (Lemma 2.2) that every Type C piece contains an "octagonal" almost normal 2-sphere. To see that each Type B piece also contains an almost normal 2-sphere, it is useful to introduce an ordering on the normal family Σ. It comes naturally by picking a vertex of T and looking at the complementary region (which is not a Type A piece) of the trivial normal 2-sphere around it. Topologically we get a 3-ball containing our normal family Σ. Now, the ordering on Σ is induced by inclusion. For example, the trivial normal 2-sphere around the vertex we removed is the unique largest element. The smallest elements in this ordering are either trivial normal 2-spheres consisting of normal triangles only or the ones bounding Type C 3-balls.

Our task is to find an almost normal 2-sphere in a piece with more than one boundary component. Pick the largest 2-sphere in its boundary. A very nice argument in [**7**] (subclaim 2.0.1.) implies that there must be an edge in T with a subarc which runs from the largest component of the boundary to some other component and whose interior is disjoint from Σ. By taking parallel copies of the two 2-spheres connected by this arc in the piece we are looking at, and tubing them together in one of the tetrahedra in the star of the edge, we obtain our almost normal 2-sphere.

All the almost normal 2-spheres we've created are separating, because we are in S^3. By picking the right complementary component in S^3 and applying Lemma 2.7, we can simplify each almost normal 2-sphere by a sequence of elementary isotopies. Since we are only using elementary isotopies going in the same direction (towards the interior of a fixed complementary component in S^3), the whole process can be realized by an embedding of $S^2 \times I$, where the top 2-sphere is the almost normal 2-sphere we started with and the bottom one is the 2-sphere coming from Lemma 2.7 (see Figures 8 and 9).

Another important point here is that the whole isotopy never leaves the Type B (or C) piece it started in. This is true simply because an analogous statement holds for each elementary isotopy. This implies that the isotoped surface coming from Lemma 2.7 will be contained in the interior of the piece containing the almost normal 2-sphere we started with.

In a Type C piece, the isotopy can go in two directions because the almost normal 2-sphere in this case contains an octagon. The surface we get, when isotoping towards the interior of the piece, will have 0 weight. This follows from the observation that we can forget about all pieces in Figure 6 if we compress each annulus with a length three normal curve in its boundary. This would then give a family of normal 2-spheres contained in the Type C piece which is a contradiction. Therefore, the 2-sphere we end up with has to miss the 1-skeleton.

Similar reasoning tells us that an isotopy in the other direction in the Type C piece has to end with a 2-sphere, intersecting the 2-skeleton T^2 in normal curves parallel to the ones coming from the boundary of the piece we are looking at and possibly in some simple closed curves which miss the 1-skeleton.

The almost normal 2-spheres in Type B pieces that we are going to consider, will never contain an octagon. We will therefore be isotoping in one direction only. Using the same kind of arguments as before, we can conclude that the 2-sphere we end up with consists of boundary components of the Type B piece (all except the two we started with, see Figure 9), tubed together by pieces depicted in Figure 6. It should be noted that if a Type B piece has only two boundary components, then the isotoped 2-sphere does not intersect the 1-skeleton.

Let Λ be the following collection of 2-spheres: In every Type C piece just take an "octagonal" almost normal 2-sphere which exists by Lemma 2.2. In each Type B piece take a copy of the almost normal 2-sphere described above with the annulus connecting two normal pieces moved by a natural isotopy, so that it envelops the edge it is parallel to. The 2-spheres from Λ in Type B pieces are therefore normal in all the tetrahedra of T, except in the ones contained in the star of the edge we isotoped over.

The sequences of elementary 3-balls, corresponding to the supports of elementary isotopies, yields the additional structure (on Type B and C pieces) that is required to define the subdivision S. Elementary discs and an element in Λ chop up each Type C piece of $S^3 - \Sigma$ (see Figure 8).

In the case of a Type B piece, the element of Λ will, after the isotopy, consist of all but two boundary normal 2-spheres tubed together by pieces described in Figure 6. Again, the Type B piece in question can be decomposed into (many) 3-balls and two punctured 3-balls. The two punctured 3-balls come from an element in Λ we started our isotopy on, and from its isotope after we've performed all elementary isotopies on the "tubed" almost normal 2-sphere (see Figure 9). The rest of the Type B piece is decomposed into 3-balls by all the elementary discs required for this isotopy.

After the isotopy from Lemma 2.7, what's left in each tetrahedron of the complementary component we isotoped into, are just 3-balls bounded by the pieces from Lemma 2.7 on one side and possibly some normal pieces of the elements in Σ on the other. Schematically, the situation after the isotopy is depicted by Figure 9.

Now we want to triangulate all of these 3-balls (the elementary ones as well as the ones that are left over in the component we were isotoping into),

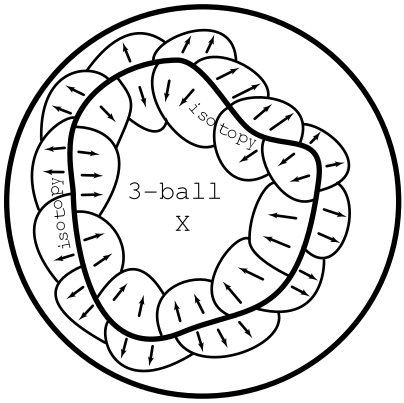

Figure 8. Two sequences of elementary isotopies in a Type C piece.

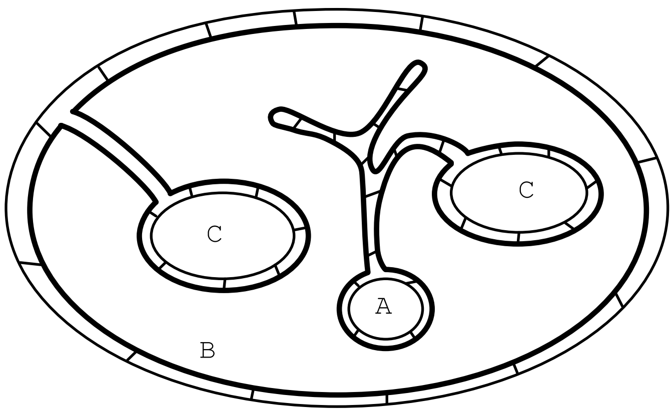

Figure 9. The "tubed" almost normal 2-sphere and the isotoped 2-sphere in a Type B piece.

in all the pieces of the complement of Σ, by simple shellable triangulations. Since all the processes described above induce polyhedral structures in the boundaries of all the 3-balls (this will be described in detail in Section 5) in question, subdividing the boundary 2-spheres in order to obtain genuine triangulations and then coning them, does the job. Doing so in every piece of the space $S^3 - \Sigma$ exhausts the whole 3-sphere and therefore completely determines the subdivision S.

The fact that all these cones are indeed shellable, is proved in [**4**] (Lemma 5.4). Here we are relying on the property that all the bounding 2-spheres we'll need to cone in the process, are triangulated by combinatorial triangulations. The reason why we want these 3-balls to be shellable is simply because each elementary isotopy can then be realized by a shelling of the corresponding 3-ball.

So what we really want from the cones on the 2-spheres above, is to be shellable without ever having to shell from the faces contained in a fixed disc, which lies in the bounding 2-sphere. This disc is just a 2-manifold along which the 3-ball, we are trying to triangulate, is glued onto the rest of the (Type B or C) piece. This can always be achieved since a cone on a disc with a combinatorial triangulation can be shelled "from the side" just by coning the shelling procedure of the disc itself.

The simplifying process works its way up the ordering of the normal 2-spheres in Σ. First we change the subdivision S in each Type C piece (which are smallest elements in our ordering), making it a cone on the unique boundary component. In Section 4 we will discuss how to implement elementary shellings from a 2-sphere boundary component by Pachner moves, if on the other side of that 2-sphere we have a cone on it. Using that construction we can pick a 3-ball piece (in some 3-simplex), contained in the 3-ball X from Figure 8, and turn the whole 3-ball X, bounded by the 2-sphere coming out of Lemma 2.7, into a cone on its boundary. This is simply because the complement (in X) of the coned 3-ball piece we picked, is shellable. That follows from the observation that all the (coned) 3-ball regions from Figure 6, the 3-ball X is made of (we already know that the first possibility in Figure 6 can not occur) can be viewed as vertices of a graph whose edges correspond to the discs in the interiors of the 2-simplices of T. Since this graph is a tree (this follows from the fact that the isotoped 2-sphere bounds a 3-ball), there is a "global" shelling strategy for the complement of the piece we picked in the 3-ball X. This can be made simplicial by shelling one cone at a time.

So now we can assume that the 3-ball X is coned. We can carry on by shelling (in the reversed order) all elementary 3-balls (and the 3-ball corresponding to the natural isotopy) involved in the isotopy taking the almost normal 2-sphere Y to the boundary of X. By this stage, we've changed the subdivision S so that it looks like a cone on Y in the Type C piece we are looking at. Above it, S is still unchanged. We can now do the same thing towards the boundary of the piece we are considering, again using the shellable nature of the subdivision S in all appropriate 3-balls.

What we have now is a cone on the isotope of Y. Now we concentrate on the $S^2 \times I$ region between the isotope of the "octagonal" almost normal 2-sphere Y and the single boundary component of the Type C piece. We first shell all the 3-balls from that region that are bounded by pieces in Figure 6. We thus obtain a cone on a normal 2-sphere which is parallel to the bounding 2-sphere of our piece. Since all the regions between any two parallel normal pieces are cones as well, we can shell them one by one and therefore get a cone on the boundary of our piece. Here we are relying on the fact that the normal structure on the bounding 2-sphere is shellable. In general this needn't hold, but the technical assumption that we are going to make on our triangulation T at the beginning of Section 5 will guarantee this property. This completes the simplification of the triangulation S in all Type C pieces.

Take a Type B piece and assume that all normal 2-spheres, strictly smaller than the largest normal 2-sphere in its boundary, already bound coned 3-balls. The strategy now is similar to the one we used in Type C pieces. Since the tube of the 2-sphere element of Λ in our piece runs from the largest boundary component to some other boundary component, we can deduce that all other normal 2-spheres in the boundary that are going to be tubed together by the isotoped 2-sphere (see Figure 9), are going to bound cones on one side.

We first shell all the regions which are bounded by two parallel normal pieces and lie between a normal boundary component and the isotoped 2-sphere. We can do that by expanding the cone structures on the other side of the boundary components of our Type B piece, that exist by assumption. Now the 3-ball bounded by the isotoped 2-sphere from Λ is again chopped up into 3-ball regions that are glued together along discs contained in the interiors of 2-simplices of T. Like before, there is a sequence of elementary shellings which gives a way of changing the triangulation of the 3-ball bounded by the isotoped 2-sphere from Λ to the cone on its boundary.

We will now mimic what we did in Type C pieces. Let's take the sequence (in the reversed order) of all elementary 3-balls coming from the elementary isotopies needed to push the 2-sphere element of Λ in our Type B piece, down to the 2-sphere which now already bounds a coned 3-ball. Using this sequence in the same way as above, we can change the triangulation S in our piece to the cone on the 2-sphere element of Λ. It is now obvious how to simplify the remains of the subdivision S in the Type B piece we are looking at. So we've managed to transform the subdivision in our piece into a cone on the largest boundary component.

Now we want to make sure that the techniques described above suffice for the total simplification of the subdivision S. Let's assume that K is a triangulation of S^3 and that S is its subdivision containing Σ as a subcomplex. Let x be the vertex of K inducing an ordering on Σ. Let's also assume the following property: If all elements of Σ, smaller then a given normal 2-sphere A in Σ, bound coned 3-balls, then using Pachner moves we can change the triangulation of the 3-ball bounded by A into a cone on A, without altering the simplicial structure of A. Then we claim that we can transform S, using Pachner moves only, into a cone on x glued to another copy of itself via an identity on the boundary.

To see this, we'll use a simple induction on the depth of elements in Σ. A normal 2-sphere in Σ is of *depth* k if it is greater than precisely k elements of Σ.

We can use our assumption for the 2-spheres of depth 0. Let A in Σ be of depth $(k+1)$ and assume we've coned all the 2-spheres of depth smaller or equal to k. Any 2-sphere smaller than A is of depth at most k. So we can use the assumption again. This proves our claim and therefore completes the simplification process.

4. Pachner moves and shellable triangulations.

In this section we are going to establish a relationship between elementary shellings and Pachner moves. We will do this in dimensions two and three. Both cases will play a crucial role in building and simplifying the subdivided triangulation S. Let's start by stating precisely what we mean by shelling.

Definition. Suppose that M' is a submanifold of a triangulated n-manifold M with boundary. If there exists an n-simplex Δ in the triangulation of M with the property that $\Delta \cap \partial M$ is a combinatorial $(n-1)$-disc, such that M' equals the closure (in M) of the complement $M - \Delta$, then we say that M' is obtained from M by an *elementary shelling*.

An elementary shelling is quite similar to an elementary collapse of the top dimensional simplex. The crucial difference lies in the fact that here we stipulate explicitly that the resulting space has to be a manifold.

Another thing which is worth mentioning is that the boundaries ∂M and $\partial M'$ differ by a single $n-1$ dimensional Pachner move.

A sequence of such elementary shellings is called a *shelling*. Saying that a triangulation of an n-manifold is *shellable* simply means that there exists a sequence of elementary shellings which will reduce the triangulation down to a single n-simplex. Since the homeomorphism type of the manifold in question does not change under an elementary shelling, it is clear that n-balls are the only candidates to have shellable triangulations. It is for example very well-known that any combinatorial triangulation of the two dimensional disc is always shellable. As it was mentioned before, Lemma 5.4 in [**4**] and the above observation about discs together imply that a cone on any combinatorial triangulation of the 2-sphere constitutes a shellable triangulation of the 3-ball.

Now we are going to express all possible elementary shellings by Pachner moves in the following three dimensional situation. Suppose we had a triangulated 3-manifold and we wanted to make an elementary shelling from a 2-sphere boundary component. Suppose further that on the other side of this 2-sphere, we had a cone on it. We have to consider three different cases according to the number of faces of the 3-simplex we are shelling, which are contained in the boundary 2-sphere.

The first case, where we have a single triangle in the boundary, is dealt with by Figure 10.

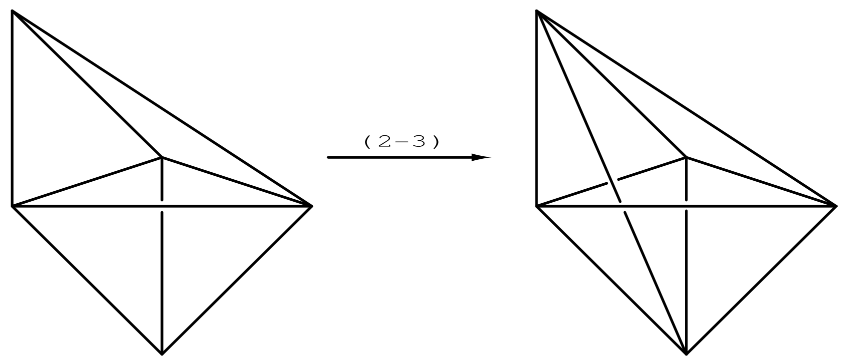

Figure 10. A single free face requires one (2-3) Pachner move.

We should note that before making the (2-3) move in Figure 10, the top 3-simplex is contained in the manifold, while the bottom one belongs to the cone. After the move, all three 3-simplices are contained in the altered cone.

The second case is the one where we have two faces in the boundary. It is clear from Figure 11, that a single (3-2) Pachner move suffices.

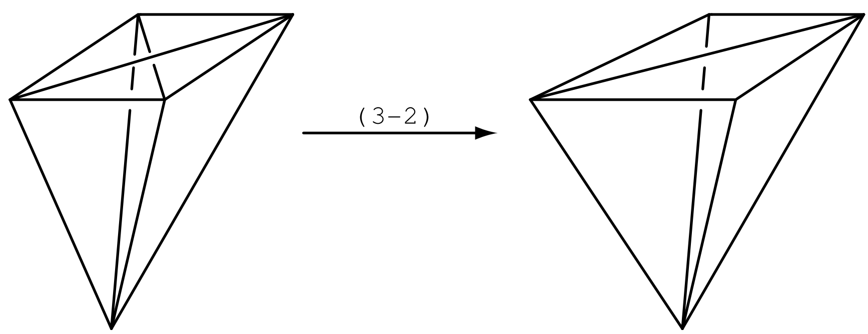

Figure 11. A single (3-2) move implements the shelling with two free faces in the boundary.

Finally, we have to deal with the situation where the 3-simplex we want to shell has three of its faces in the boundary 2-sphere. The top 3-simplex on the left of Figure 12 is the one we want to shell next, while the other three are contained in the cone. It is obvious that a single (4-1) Pachner move does the job.

Putting all these facts together, we've seen that in the setup described above, each elementary shelling corresponds to a single Pachner move. So if we want to bound the number of Pachner moves required for the simplification of the subdivision S, all we need to do is to count the number of tetrahedra in S. This will be dealt with in Section 6.

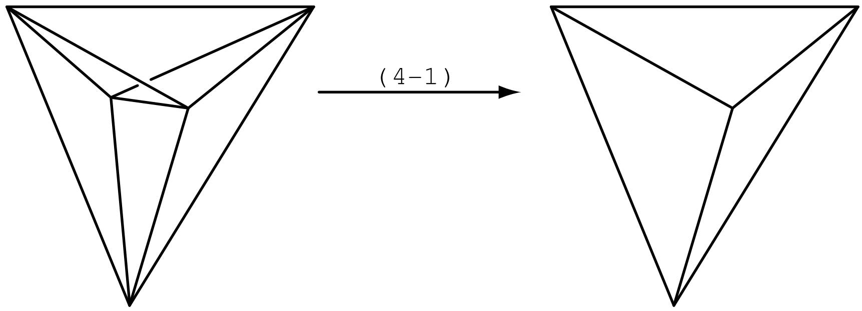

Figure 12. A single (4-1) move completes the elementary shelling.

Before we go on to discuss the two dimensional case, we need to prove the following slightly technical lemma which connects collapsing of an edge with Pachner moves. It will be of use to us in Section 5.

Lemma 4.1. *Let x be a vertex in a combinatorial triangulation of S^2 containing n 2-simplices. Assume further that the star of x is an embedded* PL *disc, triangulated by k triangles. Let e be the unique edge in the* 3*-ball, triangulated as a cone on S^2, running between x and the cone point. The triangulation of the same* 3*-ball obtained by crushing the edge e, and thus flattening its star, can be constructed by $(n-k+1)$ Pachner moves used on the original* (*coned*) *triangulation.*

Proof. The 3-ball from the lemma can be view as a union of the following two PL 3-balls: The star of the edge e and the cone on the disc in the bounding S^2, which is the complement of the star of the vertex x on the 2-sphere.

The triangulation we are aiming for is equal to the triangulation of the latter 3-ball. We therefore want to flatten the star of the edge e down to the cone on the link of e.

This can be achieved by "moving" the cone points of the 3-simplices in the second of the two 3-balls described above, from our initial cone point to the vertex x. Such a 3-simplex, having a face in S^2, which is adjacent to the star of x, can be moved by a (2-3) move or its inverse, depending on the number of edges it has in common with the star of x.

Repeating this for all (but one) 3-simplices in the cone on the disc $S^2 - \text{int}(\text{star}(x))$ almost does the job. All we have to do at this stage, is to use a single (4-1) move on what's left of the two 3-balls described above.

We should also note that the sequence of (2-3) moves and their inverses, we used to alter the initial triangulation, can always be found. This follows from the well-known fact that every combinatorial triangulation of a PL 2-disc is shellable. $\square$

The rest of this section will be devoted to two dimensional Pachner moves and their relationship with elementary shellings. In fact, what we want to do is to transform any given triangulation of a disc into a cone on its boundary, using Pachner moves only.

In dimension two there are three possible moves at our disposal. They are given by Figure 13.

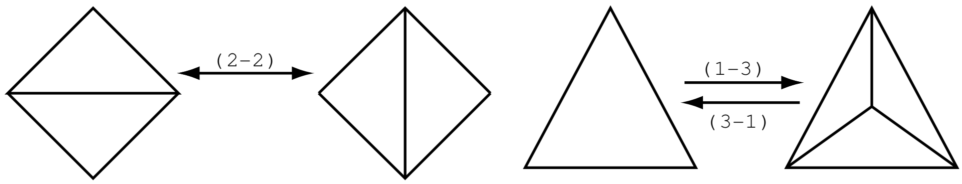

Figure 13. Two dimensional Pachner moves.

The simplifying procedure for any PL disc is described by the next lemma.

Lemma 4.2. *Any combinatorial triangulation of a piecewise linear disc with n triangles can be altered into a cone on the boundary of the same disc by n Pachner moves.*

There are two reasons why we can (and have to) assume that the triangulation of the disc is combinatorial. The first one is that in what follows, we can easily guarantee this property for all the discs we are going to be using our Lemma 4.2 on. The second one is that the proof of the above lemma relies on the fact that any triangulation of a disc is shellable, a fact not entirely correct (with our definition of an elementary shelling) if we allow for non-combinatorial triangulations.

Proof. Since the triangulation of our disc is shellable, we can index all the simplices in it by numbers from 1 to n, so that the increasing integers specify a way of reducing our triangulation down to a single triangle. The 2-simplex that's left has index n. Let's make a (1-3) move on it. The 2-simplex corresponding to $n-1$ has to share a unique edge with it. Making a (2-2) move over this edge, changes our original triangulation in the last two 2-simplices to a cone on the boundary of the disc that they compose. The rest of the triangulation is unchanged at this stage.

Noticing that the union of the last k 2-simplices in our sequence always gives a disc, makes the following induction possible. Say that we already have a cone on the boundary of the disc which is the union of the last k 2-simplices and that the rest of the triangulation we started with is unchanged. If the triangle corresponding to $n-k-1$ has a single edge in common with our cone, we act as before (a single (2-2) move suffices). If it has two faces in common, a single (3-1) move finishes the proof. $\square$

5. The subdivided triangulation.

Let T be a possibly non-combinatorial triangulation of S^3 with t tetrahedra. Let's also make the following technical assumption on T: Each edge in the 1-skeleton of T appears at most once as an edge of any 3-simplex in T. This assumption does not imply that the triangulation T is combinatorial, but it is certainly satisfied by all combinatorial triangulations of S^3. We are making it at this stage because it is going to simplify some of the processes we'll have to invoke later on. It will also become clear that any triangulation can be altered so that it has this property by linearly (in t) many Pachner moves.

In this section, we shall describe the subdivision S of the triangulation T and also bound the number of Pachner moves required to construct it.

Let Γ be the union of all discs needed to perform all elementary and natural isotopies in all the pieces of $S^3 - \Sigma$. We should note that the number of elements of Γ, coming from a single 2-sphere in Λ, is bounded above by the number of times the almost normal sphere in question intersects the 1-skeleton. An explicit bound on the number of elements of Γ will be given later.

An elementary disc from Γ will intersect every tetrahedron in the star of the edge we are isotoping over in a disc region (see Figure 14).

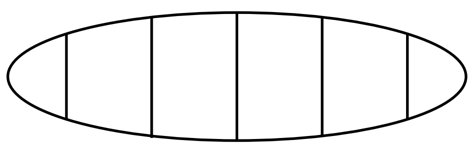

Figure 14. Regions, in a disc in Γ, correspond to tetrahedra in the star of an edge.

In each tetrahedron, the operation of adding in this disc will consist of gluing in a length four (respectively two) disc so that two (respectively one) arcs in its boundary are contained in the isotoped almost normal surface, and the other two (respectively one) lie in the boundary of the tetrahedron.

We are now in the position to describe the subdivision of the polyhedron

$$T^2 \cup \Sigma \cup \Lambda \cup \Gamma$$

which will be a subcomplex of the triangulation S. In fact, the simplicial structure of the polyhedron $T^2 \cup \Sigma \cup \Lambda \cup \Gamma$ will play a crucial role in the simplifying process and will also be of significance in the definition of the subdivision S.

All the normal 2-spheres in Σ will inherit the PL structure from their normal structure. The normal triangles in Σ will become 2-simplices, while the normal quadrilaterals will be subdivided into two 2-simplices by a diagonal.

The PL structure of the almost normal 2-spheres in Λ will be a subdivision of the normal and almost normal pieces. We will subdivide them according to the markings on them, made by discs in Γ, where we define a *marking* on a normal or an almost normal piece to be an arc of intersection of the piece with a disc in Γ. We know, that each element of Γ is chopped up into discs of lengths two or four in each tetrahedron (as in Figure 14).

The disc regions of the elements of Γ of length four can only leave a marking on a normal piece of an almost normal 2-sphere going from one normal arc to another. There can only be three such markings on a triangle and four of them on a quadrilateral, one for each corner. On an almost normal octagon, immediately after we glue on our first disc corresponding to the natural isotopy, we end up with two triangles. So there can be at most six markings on an octagon, coming from the discs of length four.

The discs of length two will either leave a marking running from a normal arc to some other marking or simply running between two markings. Because each marking is parallel to some edge of the normal piece that it lies on and because we can not get more than one marking of the same kind, superimposing all the possible markings on normal and almost normal pieces is described by Figure 15.

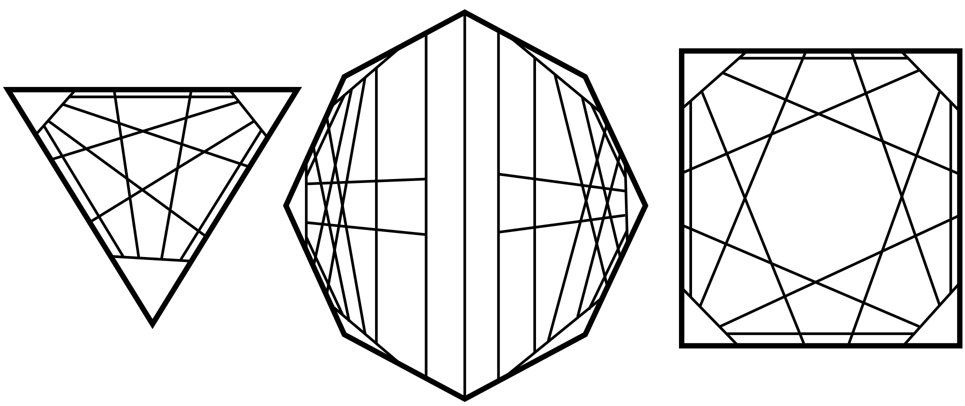

Figure 15. The polyhedral structure of normal and almost normal pieces of elements in Λ.

The almost normal piece which is obtained by tubing together two normal pieces can be treated in the same way, since we could view it as an annulus around an edge between two normal 2-spheres. This annulus consists of discs of length four in each tetrahedron in the star of the connecting edge. These discs will be glued on the pairs of normal pieces yielding non-normal pieces, similar to the ones we get during the isotopy of the surface F from Lemma 2.7. The PL structure on such a piece will come from the PL structure on the two parts of normal pieces it consists of, and from the

PL structure on the glued in disc, which we haven't yet described. These glued in discs from the annulus behave in the same way as the disc regions from elements in Γ (see Figure 14).

In the next paragraph we shall see that each of these disc regions can be triangulated by at most 6 triangles. Counting the regions in the normal and almost normal pieces in Figure 15 and triangulating each region (if it is not a triangle already) by coning from one of the vertices in its boundary, we can see that each piece, including the ones coming from the "tubed" almost normal 2-spheres, contains less than 200 2-simplices. We should also note that the described subdivision of the pieces is combinatorial.

Now, we have to put a PL structure on the elements of Γ. We've noted before (Figure 14) that each elementary disc in Γ consists of disc regions of lengths two or four. Once we've glued in a disc from Γ, the disc regions in it give us a polyhedral structure on it. Further gluings will however subdivide this structure. Concentrating on a single disc region A of our element in Γ, we note that all further gluings of disc regions of length four will miss A completely and therefore not change it at all. Disc regions of length two can add in a further arc on A which runs parallel to the arc(s) in its boundary, contained in the 2-skeleton T^2. Since this can only happen once per boundary arc of A in the 2-skeleton T^2, we can add at most two arcs in each disc region of any element in Γ. So a disc in Γ will in the end look exactly like the disc in Figure 14 with less than $3t$ disc regions. This follows from the assumption we made at the beginning of this section, since it implies that a star of an edge can contain at most t tetrahedra.

The arcs in the boundaries of disc regions of elements in Γ that leave markings on normal and almost normal pieces of the elements in Λ will be subdivided further by the vertices coming from the points of intersection of the markings (see Figure 15). An arc in the boundary of the length two disc region (i.e., the one that's leftmost or rightmost in Figure 14) will get at most 16 vertices in this way, while an arc in the boundary of the length four disc region will contain at most 5 such vertices (see Figure 15).

All these observations about the polyhedral structure of the discs in Γ imply that each disc region corresponding to a single tetrahedron in T, will be triangulated by no more than 20 triangles. So we can triangulate any element from Γ by less than $20t$ triangles. Again, the triangulation we get is combinatorial.

Finally, we need to induce a PL structure on the 2-skeleton T^2. Normal and almost normal simple closed curves bounding pieces of elements of Σ and Λ will partition the 2-skeleton T^2 into piecewise linear regions and thus induce a polyhedral structure on it. We only have five nontrivial complementary regions in the boundary of every tetrahedron in T. They are as in Figure 16.

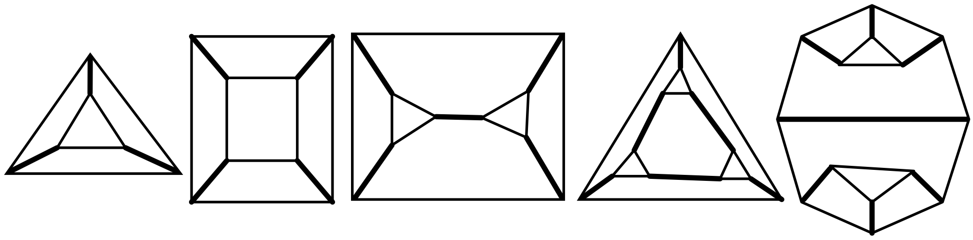

Figure 16. Regions in the boundary of a 3-simplex bounded by normal and almost normal simple closed curves.

So topologically we have two annuli, two twice punctured discs and one three times punctured disc. The technical assumption on the triangulation T, we made at the beginning of this section, implies that all these surfaces are embedded in the 3-sphere.

We also need to take into account the discs in Γ which will subdivide further the polyhedral structure that the surfaces in Figure 16 already have. Each disc region of an element in Γ will give a further arc in one of the regions in Figure 16. This arc will run from one normal arc in the boundary of the region to the other. Its end points are vertices of the subdivision of normal and almost normal pieces of the 2-sphere elements from Λ. It is worth noting that an arc in the boundary of a tetrahedron, coming from a disc in the family Γ, will neither connect two segments in the 1-skeleton T^1 nor will it connect a normal arc with a segment in T^1. Since a normal arc can have at most 4 vertices in its interior (Figure 15), it follows that we will never have to add in more than 50 arcs per planar surface (adding up all the possibilities in all the regions in Figure 16) in T^2.

We can now obtain the simplicial structure on the 2-skeleton just by coning from one of the vertices of each disc subregion of the planar surfaces in Figure 16. It now follows that each surface in Figure 16 is triangulated by less than 200 2-simplices.

Now we are in the position to describe completely the subdivision S of the triangulation T we started with. As it was said before, the polyhedron $T^2 \cup \Sigma \cup \Lambda \cup \Gamma$ with its simplicial structure is going to be a subcomplex of S. Lemma 2.7 tells us that the complement of the polyhedron $T^2 \cup \Sigma \cup \Lambda \cup \Gamma$ in each tetrahedron of T is just a union of 3-balls. The boundary 2-spheres of these 3-balls are embedded by the assumption we made at the very beginning of this section. They also inherit a PL structure from $T^2 \cup \Sigma \cup \Lambda \cup \Gamma$ which is combinatorial. Every complementary 3-ball can thus be triangulated by adding a vertex in its interior and coning its boundary. Since these 3-balls exhaust the whole 3-sphere, the cones completely determine the subdivision S. We should also note that all these coned 3-balls are in fact shellable because their bounding 2-spheres are triangulated in a combinatorial fashion.

The rest of this section will be devoted to obtaining the subdivision S from the triangulation T using Pachner moves. The basic tool for achieving this end will be the procedure called *changing of cones.*

Suppose we had two PL discs D and E with isomorphic simplicial structure on their boundaries. Let the union $D \cup E$ denote the PL 2-sphere obtained by gluing the two discs together via a simplicial isomorphism on their boundaries. What we want is an algorithm to transform the cone on D, denoted by CD, to the union of cones $CE \cup C(D \cup E)$, without changing the triangulation of D. This is described schematically by Figure 17.

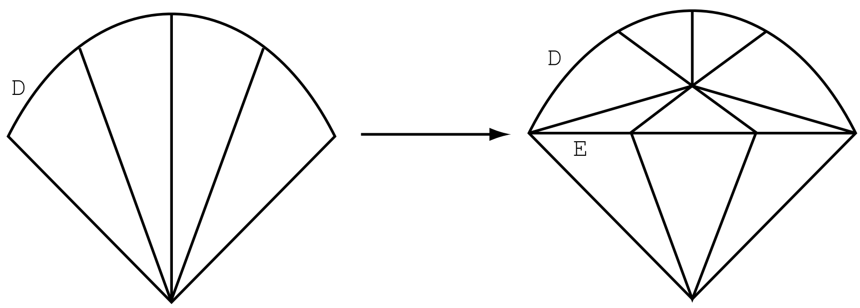

Figure 17. The changing of cones.

We have the following lemma giving a bound on the number of Pachner moves required for changing of cones.

Lemma 5.1. *Let discs D and E be as above, where n is the number of 2-simplices in D and m is the number of 2-simplices in E. Then we can perform the changing of cones using less than $4(n+m)$ Pachner moves.*

Proof. We will divide the process into three steps. First, we glue a cone on the cone on the boundary of D onto the bottom part of the boundary of CD (Figure 18). This is a reversed process to destroying an edge which connects the two cone points of the bit that we glued on. It can therefore, by Lemma 4.1, be accomplished by less than $(n+1)$ Pachner moves.

In the second step we perform the same move again, i.e., we glue the cone $C(C(\partial D))$ onto the space we've got so far (Figure 18). This again requires not more than $(n+1)$ Pachner moves.

The space we've created can be described as a suspension of $C(\partial D)$ glued onto the cone on the disc D. We know that we can transform the cone triangulation of the disc $C(\partial D)$ into the triangulation of E by using not more than $(n+m)$ two dimensional Pachner moves (Lemma 4.2). It is also

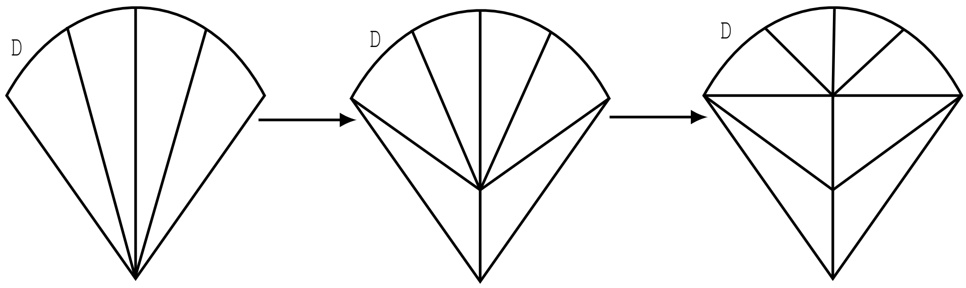

Figure 18. The first two steps in the changing of cones.

clear that in the suspension setting, each (1-3) move (or its inverse) can be realized by one (1-4) and one (2-3) move. A (2-2) Pachner move can be realized by a (2-3) and a (3-2) move. Putting all this together implies our bound. □

The changing of cones will help us produce all the necessary cones in the triangulation S. Now we have at our disposal all the tools required, to bound the number of Pachner moves needed for obtaining the subdivision S from the triangulation T.

The whole process will be divided into five stages. We'll start by describing each one of them, and then we'll bound the number of moves we made.

1) Add a vertex into every tetrahedron and every triangle of the triangulation T and cone.
2) Subdivide the 1-skeleton of T to get a subcomplex of S, and keep the triangulation in the 3-simplices of T coned.
3) Subdivide the 2-skeleton of T to get a subcomplex of S, and keep the triangulation in the 3-simplices of T coned.
4) Chop up tetrahedra of T by the appropriate normal and almost normal pieces and triangulate the complementary regions by coning them from a point in their interior.
5) Chop up the complementary regions of 4 by length two and length four disc regions of elements in Γ. Cone the complements.

We note that Step 3 can be accomplished by suspending the process in Lemma 4.2. Steps 4 and 5 are possible by Lemma 5.1.

Adding a vertex into each 3-simplex in T takes t(1-4) moves. Adding one into a triangle of T takes two Pachner moves: One (1-4) move followed by a (2-3) move. So Step 1 amounts to $5t$ Pachner moves since there are precisely $2t$ triangles in the triangulation T.

We should note that the subdivision we get after Step 1 will always satisfy the technical condition we stipulated at the beginning of this section. This is simply because every tetrahedron of this subdivision contains precisely one edge from the 1-skeleton of T. Its other edges are embedded in the 2-simplices and in the tetrahedra of T. It is also clear that this subdivision contains $12t$ 3-simplices. So the worst case scenario would make us do Step 1 at the very beginning and then do the simplification process (that we've been describing) on that subdivision. So once we work out the bound for this simplification procedure, we have to substitute each t in the formula with $12t$.

Let's go back to the construction of the subdivision S. First we want to bound the number of vertices of S in each edge of the triangulation T. By Lemma 2.6 it follows that there are at most $3 \cdot 2^{110t^2}$ normal arcs in any triangle of T, coming from all elements in Σ and Λ. Since each normal arc contributes at most one point of intersection with a single edge, we will have less than $3 \cdot 2^{110t^2}$ vertices on any edge in the 1-skeleton T^1. Since there are less than $5t$ edges all together (an Euler characteristic count), the number of vertices of the triangulation S, contained in T^1 will be bounded by $15t2^{110t^2}$.

The star of any edge in T contains at most $2t$ 3-simplices in the subdivision we have so far. Creating a vertex on this edge can obviously be done in the following way: First make a (1-4) move on one of the simplices in the star of the edge. Then do a sequence, of length at most $2t-2$, of (2-3) Pachner moves. Now the addition of the vertex can be finished off by a single (3-2) Pachner move. All together this procedure takes not more than $2t$ Pachner moves. Step 2 will thus require no more than

$$30t^2 2^{110t^2}$$

Pachner moves.

We already know that there will be at most $3 \cdot 2^{110t^2}$ normal arcs in any triangle of T. So the number of regions in a 2-simplex in the 2-skeleton T^2 is therefore bounded by the same number (plus one). These regions correspond to the regions in the surfaces from Figure 16 and will thus be triangulated by less than 20 2-simplices. So any triangle in T will be subdivided by at most $60 \cdot 2^{110t^2}$ 2-simplices. By Lemma 4.2, this configuration can be obtained by $60 \cdot 2^{110t^2}$ two dimensional Pachner moves (we should notice here that before starting the process from Lemma 4.2, the triangles of T were subdivided as cones on their boundaries). Suspending this process and doing it for all $2t$ 2-simplices in T^2 yields an upper bound of

$$3 \cdot 10^2 t 2^{110t^2}$$

Pachner moves used in Step 3. This is because every two-dimensional Pachner move requires 2 three-dimensional ones.

The number of 3-ball regions, the elements of Σ and Λ produce in all tetrahedra of T, is equal to the number of normal and almost normal pieces in all the 2-spheres from Σ and Λ (plus t). So it is bounded above by

$3 \cdot 2^{110t^2}$. Using Lemma 5.1, we are going to change the cone structure in every tetrahedron in T. This will be accomplished, step by step, starting from the vertices of the tetrahedron and moving towards the cone point in its interior. At each stage we have to change a disc consisting of one of the surfaces in Figure 16, where all but one of its boundary components already have their corresponding normal and almost normal pieces glued in (that makes it a disc), to a disc coming from the only normal or almost normal piece that hasn't yet been introduced. Since we want the region between the two discs we've just described, to be coned, Lemma 5.1 is precisely what is needed. It is also obvious that the disc D from Lemma 5.1 will in this situation never contain more than 800 triangles (this follows from the counts we did when defining the subdivision S), while the disc E, which is just a normal or an almost normal piece, will be triangulated by less than 200 2-simplices. So in a single 3-ball region, we'll make less than $4 \cdot (800 + 200)$ Pachner moves (Lemma 5.1). In order to complete Step 4 in all the tetrahedra of T, we need to make

$$12 \cdot 10^3 2^{110t^2}$$

Pachner moves.

The number of discs in Γ, coming from a single element in Λ, is bounded above by half the number of times the 2-sphere in question intersects the 1-skeleton. We already know that there are at most $3 \cdot 2^{110t^2}$ vertices on any edge in the 1-skeleton of the triangulation T. Since there are less then $5t$ edges in T^1, the number of elements in Γ is bounded above by $\frac{1}{2} 15t 2^{110t^2} < 10t 2^{110t^2}$.

Each of the discs in Γ has at most t disc regions (by the assumption from the beginning of this section), coming from the 3-simplices in the star of the edge the particular disc corresponds to. Each disc region is triangulated by strictly less than 20 triangles. A disc region in an element of Γ will correspond to the disc E in Lemma 5.1.

The boundary of each disc region is a subcomplex in the boundary of a coned 3-ball. One of the complementary discs bounded by this simple closed curve, in the boundary of the coned 3-ball, will correspond to the disc D in Lemma 5.1. In the case of a disc region in an element of Γ having two arcs in its boundary embedded in the 2-skeleton T^2, the disc corresponding to D we were discussing before will contain six 2-simplices (two in normal or almost normal pieces and four in the 2-skeleton T^2).

Let's look at the case of a disc region from an element in Γ that intersects the 2-skeleton of T in a single arc (i.e., the leftmost or the rightmost region in Figure 14) and corresponds to an elementary isotopy. The number of triangles of the complementary region (in the bounding 2-sphere) we are interested in will then be smaller than the sum of the numbers of 2-simplices in the following surfaces: The disc in the 2-simplex of T our disc region is parallel to, the disc in the 2-simplex of T containing a bounding arc of the disc region we are gluing in, regions in at most three normal triangles or regions in a normal triangle and a normal quadrilateral or regions in two normal quadrilaterals, at most two discs contained in two distinct regions in the elements of Γ. Bounds for the numbers of 2-simplices for the above surfaces are as follows: 20, 2, $3 \cdot 30$ or $2 \cdot 30$ or $2 \cdot 30$, $2 \cdot 2$ respectively. What happens with the disc regions belonging to the elements of Γ that come from natural isotopies? In that case the disc D from Lemma 5.1 is composed of the following surfaces: Roughly a half of an almost normal octagon, three discs contained in the 2-simplices of T, a single normal triangle. The explicit bounds in this case are: 70, $3 \cdot 20$, 70.

An upper bound on the sum of the numbers of triangles in D and E will therefore always be strictly less than 300 (we already know that a disc region in an element from Γ contains no more then 20 2-simplices). So by Lemma 5.1, we can produce our disc region in this 3-ball by less than $4 \cdot 300$ Pachner moves. All together, we have to make less than

$$12 \cdot 10^3 t^2 2^{110t^2}$$

Pachner moves in order to complete Step 5.

Summing everything up, estimating the resulting expression and substituting t with $12t$ to account for the technical assumption we made at the beginning of this section, we get the following proposition.

Proposition 5.2. *Let T be any triangulation of the 3-sphere and let t be the number of tetrahedra in it. Then the subdivision S, described at the beginning of this section, can be obtained from T by making less than $ct^2 2^{dt^2}$ Pachner moves, where the constant c is bounded above by $5 \cdot 10^6$ and the constant d is smaller than $2 \cdot 10^4$.*

6. Conclusion of the proof.

Now, we are in the position to bound the number of Pachner moves needed to simplify any given triangulation T of the 3-sphere, down to the canonical triangulation with only two tetrahedra. We will apply the shelling techniques, developed in Section 4, to the subdivision S of the triangulation T, described in Section 5.

The basic question we have to answer at this point is how many tetrahedra do we have to shell in the simplifying process. Then we can estimate the number of Pachner moves needed for the process, using the fact that each elementary shelling corresponds to a single Pachner move.

Let's bound first the total number of tetrahedra of S.

This will be accomplished in two steps. First we count the number of 3-ball regions we coned, while constructing the subdivision S, in all the tetrahedra of the triangulation T. The second step consists of bounding the number of triangles in each of the boundaries of the 3-balls mentioned above. Multiplying these two numbers gives our bound.

Lemma 2.6 implies that there are at most $3 \cdot 2^{110t^2}$ normal and almost normal pieces in all 3-simplices of T, coming from all normal and almost normal 2-sphere in $\Sigma \cup \Lambda$. We know that each piece contains at most 200 triangles. Each planar surface in the boundary of the tetrahedron (see Figure 16) contains at most 50 arcs and is triangulated by at most 200 triangles. Each 3-ball component of the complement of $\Sigma \cup \Lambda$ in our tetrahedron will thus contain less than 50 disc regions, coming from elements in Γ.

So in all 3-simplices of T we'll have not more than $50 \cdot 3 \cdot 2^{110t^2}$ 3-ball regions. Since each disc region in any element in Γ contains less than 20 triangles, 1000 is surely an upper bound on the number of triangles in the boundary of any of the 3-ball regions. There will therefore be at most $15 \cdot 10^4 2^{110t^2}$ tetrahedra in S.

Combining Proposition 5.2 and the assumption that there are precisely $12t$ tetrahedra in the triangulation T, concludes the proof of the main theorem.

Acknowledgements. I would like to thank my research supervisor Marc Lackenby for many helpful (and enjoyable) conversations about math and about everything else.

References

[1] A. Casson, *Lecture given at the topology seminar*, Berkeley, 1998.

[2] J. Hass and J.C. Lagarias, *The number of Reidemeister moves needed for unknotting.* Preprint, 1998.

[3] S.A. King, *The size of triangulations supporting a given link.* Preprint, 2000.

[4] W.B.R. Lickorish, *Simplicial moves on complexes and manifolds*, Geometry Topology Monographs, **2**, Geometry and Topology Publications, Warwick, 1999, 299-320, MR 2000j:57057, Zbl 0963.57013.

[5] U. Pachner, PL *homeomorphic manifolds are equivalent by elementary shellings*, Europ. J. Combinatorics, **12** (1991), 129-145, MR 92d:52040, Zbl 0729.52003.

[6] M. Stocking, *Almost normal surfaces in* 3*-manifolds*, Trans. Amer. Math. Soc., **352**(**1**) (1999), 171-207, MR 2000c:57045, Zbl 0933.57016.

[7] A. Thompson, *Thin position and the recognition problem for* S^3, Mathematical Research Letters, **1** (1994), 613-630, MR 95k:57015, Zbl 0849.57009.

Received January 30, 2001 and revised October 27, 2001.

Department of Pure Mathematics and Mathematical Statistics
Center for Mathematical Sciences
University of Cambridge
Wilberforce Road
Cambridge, CB3 0WB
UK
E-mail address: a.mijatovic@dpmms.cam.ac.uk