Efficient and Accurate B-rep Generation of Low Degree Sculptured Solids using Exact Arithmetic *

John Keyser Shankar Krishnan Dinesh Manocha

Department of Computer Science
University of North Carolina
Chapel Hill, NC 27599
USA
{keyser,krishnas,manocha}@cs.unc.edu
http://www.cs.unc.edu/~geom/geom.html

Abstract

We present efficient representations and algorithms for exact boundary computation on low degree sculptured CSG solids using exact arithmetic. Most of the previous work using exact arithmetic has been restricted to polyhedral models. In this paper, we generalize it to higher order objects, whose boundaries are composed of rational parametric surfaces. The use of exact arithmetic and representation guarantees that a geometric algorithm is numerically accurate and is likely to be required for perturbation techniques which handle degeneracies. We present efficient algorithms for computing the intersection curves of trimmed parametric surfaces, decomposing them into multiple components for efficient point location queries inside the trimmed regions, and computing the boundary of the resulting solid using topological information and component classification tests. We also employ a number of previously developed algorithms like algebraic curve classification, multivariate Sturm sequences, and multivariate resultants. We have implemented key parts of these algorithms and preliminary implementations indicate the performance of our algorithm to be about one order of magnitude slower than similar algorithms using IEEE floating-point arithmetic.

1 Introduction

Constructive Solid Geometry (CSG) and Boundary Representations (B-rep) are two major approaches for representing solids [Bra75, RV85, Hof89, Man88]. CSG implicitly represents a solid as an algebraic expression, and B-rep explicitly stores an object as a set of surfaces. Both these representations have different inherent strengths and weaknesses, and for most applications both are desired.

Many of the current solid modeling systems are based on B-reps, and Boolean combinations (union, difference, intersection) are some of the common operations performed. Computing the B-rep of the resulting solid (after performing Boolean operations) is an important operation in these systems. In this paper the objects correspond to sculptured solids, whose boundary can be represented using trimmed rational parametric surfaces. This is a wide family of objects and can exactly represent quadrics, tori and free-form solids.

The first systematic study of CSG to B-rep conversion appeared in [RV85] and nowadays the algorithms for conversion are relatively well understood [Hof89, Man88, CB89, Sar83, KM96]. However, the problem of *robust* and *accurate* computation of the boundary is considered one of the difficult problems in geometric and solid modeling [Hof96, Hea95, For96]. It is important that the computed B-rep be accurate, or at least topologically consistent, and this can be jeopardized by even small amounts of error in the representation of the model or in finite-precision computations (e.g. round-off errors).

A number of approaches, mostly restricted to polyhedral modelers, have been proposed for robust and accurate B-rep computation. One of the most common approaches is based on using *tolerances* with floating-point arithmetic [Jac95], however it is hard to decide a global tolerance value for all computations. To circumvent these problems, combinations of symbolic reasoning [HHK89] and adaptive tolerances [Seg90] have been proposed. Other algorithms include those based on redundancy elimination [FBZ93].

B-rep computation algorithms involve accurate evaluation of the sign of arithmetic expressions, which can present problems for floating-point arithmetic when the value of the expression is close to zero. If this problem is not properly addressed, the resulting algorithm becomes *un-*

*Supported in part by an Alfred P. Sloan Foundation Fellowship, ARO Contract DAAH04-96-1-0257, NSF Grant CCR-9319957, NSF Grant CCR-9625217, ONR Contract N00014-94-1-0738, ARPA Contract DABT63-93-C-0048 and NSF/ARPA Center for Computer Graphics and Scientific Visualization

reliable. Many algorithms based on *exact arithmetic* have been proposed for reliable numeric computation for polyhedra [SI89, For95, BMP94, Hof89]. These algorithms use a fixed upper bound on the bit-length of arithmetic required to evaluate geometric predicates. In particular, Fortune has presented an efficient algorithm based on exact arithmetic which has a small performance overhead as compared to a floating-point based implementation [For95]. Besides reliable computation, exact arithmetic allows the use of symbolic perturbation to handle degeneracies [Yap90]. The perturbation scheme greatly simplifies the implementation of the solid modeler.

There is relatively little work on robust or accurate B-rep computation algorithms for curved primitives. Algorithms to handle degenerate intersections between quadrics have been presented in [MG91]. For arbitrary degree sculptured solids, it is difficult to compute *tight bounds* on the error generated due to floating-point arithmetic. As a result, it is hard to extend algorithms based on tolerances to curved models. Furthermore, exact arithmetic for curved domains is perceived, for a number of reasons, to be *extremely slow* and *complex*. Exact arithmetic involves computations on algebraic numbers and most of the current implementations of such arithmetic (e.g. those available as part of computer algebra systems) are extremely slow. Techniques using bit-length estimates may, in the worst case, require bit-lengths which are exponential with respect to the degree of the algebraic functions [Can88, Yu92]. Moreover, many representations and predicates that are well-understood in the linear domain become rather hard in the curved domain. Overall, no good solutions are known for efficient and robust B-rep computation on curved solids.

Main Contribution: We present efficient representations and algorithms for exact boundary computation on Boolean combinations of sculptured solids. Our contributions include:

- **Representation:** We present efficient and exact representations for points, edges and surfaces using algebraic sets along with a topological representation.
- **Kernel Routines:** We identify lower-level routines where the algorithms based on floating point arithmetic are susceptible to failure. These include sign-evaluation of geometric predicates, orientation of points with respect to curves, and component classification. We present fast algorithms to perform such tests *reliably* using exact arithmetic. We refer to the resulting set of routines as *kernel routines*. The efficiency and reliability of the overall algorithm is governed by these routines.
- **B-rep Computation:** Given our kernel routines, we present an algorithm for B-rep computation.
- **Handling Degeneracies:** We identify most cases where degeneracies can affect our algorithm, and propose ways to identify and resolve some of them.
- **Implementation and Performance:** We describe the performance of a preliminary implementation of our algorithm.

The resulting algorithm and system work well on low-degree solids (composed of polyhedra, quadrics, tori, low-degree solids of revolution). In practice, most of the curved primitives of solid modeling systems are indeed low-degree. As compared to algorithms implemented in floating point arithmetic, our algorithm performs slightly more than one order of magnitude slower on low degree solids on average.

Organization: The rest of this paper is organized as follows. Section 2 discusses background material, including our representation for solids. Section 3 gives an overview of our algorithm and discusses the kernel routines which form its basis. Section 4 discusses how each of the major steps are performed. In Section 5, we present some analysis of our approach along with some performance results and an illustrating example. Section 6 discusses degeneracies and Section 7 concludes with a mention of possible areas for extensions and future work.

2 Background Material

In this section, we present our representation for a solid. Our algorithms assume that solids are specified in this format, and the B-rep of resulting solids is given in this format. We also present some background material that we use to compute the B-rep. This includes a number of algorithms from computational algebra. In particular, we shall briefly discuss our representation of algebraic numbers, techniques for root isolation using multivariate Sturm sequences, and multipolynomial resultant computation.

2.1 Representation of Solids

Every solid is represented as a set of *trimmed* parametric surface patches which define the solid boundary. We represent each surface patch $\mathbf{F}(s,t)$ as a *rational function* with *rational* coefficients. This kind of parametrization is possible for all quadric surfaces such as spheres and cylinders, surfaces of revolution, and tori. The domain of the patch is the unit square in the (s,t)-plane $(0 \leq s,t \leq 1)$. If we are given a different rectangular domain, we can always reparameterize to $(0 \leq s,t \leq 1)$.

Assumptions: *Topological* information of the solid is maintained in terms of an adjacency graph. It is similar to the winged-edge data structure [Hof89]. To start with, we assume that each of the input objects has *manifold* boundaries, and the Boolean operation is *regularized*. While it is possible to generate non-manifold objects from regularized Booleans on manifold solids, we assume for the sake of simplicity that this does not occur. It is a well-known fact that, while dealing with topological representation of curved objects, global resolution of edge ambiguities cannot be guaranteed at times [Hof89]. Some of these issues are addressed in Section 6. Given these assumptions, it can

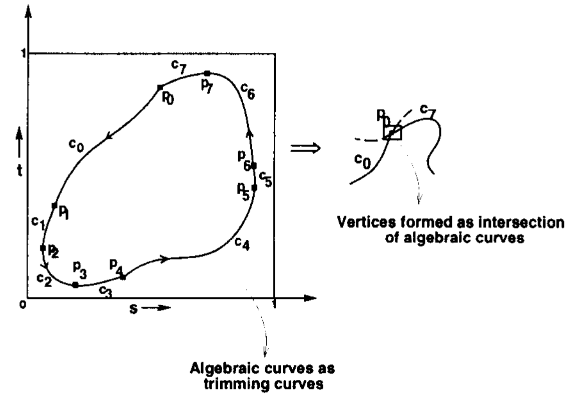

Figure 1: Representation of a trimmed patch as algebraic curve segments

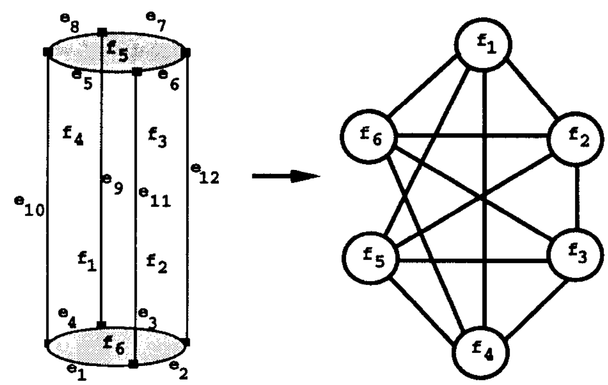

Figure 2: A cylinder and its face connectivity structure

be shown that an unambiguous topological representation is possible for a solid.

A trimmed patch consists of a sequence of curves defined in the domain of the patch such that they form a closed curve (c_i's in Fig. 1). Each c_i is a segment of an algebraic curve. The portion of the patch that lies in the interior of this closed curve is retained. Most of these trimming curves correspond to intersection curves between two surfaces. Therefore, these curves are typically algebraic curves that do not admit a rational parametrization. We represent these curve segments (c_i) by their algebraic equation and the two endpoints (p_i and p_{i+1}). The endpoints are computed by solving a set of polynomial equations, and are actually algebraic numbers (see Fig. 1). Exact representation of these numbers is discussed later in this section.

This representation of a solid lends itself to a description in terms of *faces, edges, and vertices* analogous to the polyhedral case. Each *face* is a trimmed patch. Each of the trimming curves form an *edge*, and are formed as an intersection of two surfaces (faces). Finally, endpoints of edges form the *vertices*. They can be represented as an intersection of three surfaces. Fig. 2 shows an example solid and the face connectivity structure that we maintain. We also maintain the two faces that are adjacent to each edge, and an anticlockwise order of faces around each vertex.

Representation of algebraic numbers: It was mentioned earlier that each of the vertices in the solid is defined as the intersection of three surfaces, i.e. a root of a set of polynomial equations with rational coefficients. Because of the rational parametrization of the surface, each of these equations is either univariate or bivariate. A vertex in the patch domain is therefore the common solution of two equations, $f(s,t) = 0$ and $g(s,t) = 0$. These are usually algebraic numbers, and cannot be represented exactly as finite precision numbers. Notice that a real algebraic number is the solution of an equation, $f(s) = 0$, within some interval, $a \leq s < b$. In our algorithm, we represent each algebraic coordinate as an arbitrarily small rational rectangle (i.e. an axis-aligned rectangle whose four vertices have rational coordinates). The rational rectangle is guaranteed to isolate each common root of $f(s,t)$ and $g(s,t)$ (taking into account the multiplicities of roots). The root isolation algorithm uses *multivariate Sturm sequences* as proposed by Milne [Mil92].

2.2 Multipolynomial Resultants

Elimination theory investigates the conditions under which sets of polynomials have common roots. Usually, it concerns itself with sets of n homogeneous polynomials in n unknowns, and finds the relationship between the coefficients of the polynomials which can be used to determine whether the polynomials have a non-trivial common solution.

Definition 1 *[Sal85] A* **resultant** *of a set of polynomials is an expression involving the coefficients of the polynomials such that the vanishing of the resultant is a necessary and sufficient condition for the set of polynomials to have a common non-trivial root.*

[Mac02] provided a general method for eliminating n variables from n homogeneous polynomials. The resultant is expressed as a ratio of two determinants. However, a single determinant formulation exists for $n = 2$ and 3 [Sal85, Dix08]. For $n = 3$, however, [Dix08] gives the resultant only if the three equations have the same degree. In our application, it is sufficient to compute resultants for the cases when $n = 2$ and 3. Sylvester's method [Sal85] can be used to express the resultant of two polynomials of degree m and n respectively as a determinant of a matrix with $(m+n)$ rows and columns. For the polynomials,

$$f^n(x) = a_n x^n + a_{n-1} x^{n-1} + \ldots + a_1 x + a_0 \qquad (1)$$

and

$$g^m(x) = b_m x^m + b_{m-1} x^{m-1} + \ldots + b_1 x + b_0 \qquad (2)$$

where $n \geq m$, the Sylvester's resultant is

$$\begin{vmatrix} a_n & a_{n-1} & \dots & & a_0 & 0 & \dots & 0 \\ 0 & a_n & a_{n-1} & \dots & a_0 & 0 & \dots & 0 \\ 0 & \dots & 0 & a_n & a_{n-1} & & \dots & a_0 \\ 0 & \dots & & 0 & b_m & b_{m-1} & \dots & b_0 \\ 0 & \dots & 0 & b_m & b_{m-1} & \dots & b_0 & 0 \\ b_m & b_{m-1} & \dots & b_0 & 0 & \dots & & 0 \end{vmatrix} \quad (3)$$

The problem of computing the implicit representation of a parametric surface $\mathbf{F}(s,t) = (X(s,t), Y(s,t), Z(s,t), W(s,t))$ involves eliminating s and t from the three polynomials

$$X(s,t) - xW(s,t) = Y(s,t) - yW(s,t) = Z(s,t) - zW(s,t) = 0.$$

We use Dixon's resultant [Dix08] to compute the implicit form as described in [Sed83].

Resultant computation: We use an algorithm based on multivariate interpolation [MC93] to compute the resultant of a set of polynomials efficiently. The main bottleneck in most resultant algorithms is the symbolic expansion of determinants. Most of the computer algebra systems use symbolic algorithms like polynomial manipulations for resultants, which are very expensive. Further, the magnitude of intermediate expressions grows quickly, and the memory requirements are high. The algorithm in [MC93] performs all computations over finite fields, and uses a probabilistic algorithm based on the Chinese Remainder Theorem to recover actual coefficients.

2.3 Multivariate Sturm sequences

Here, we describe briefly the algorithm proposed by Milne [Mil92] to compute the number of common real solutions of n polynomials in n variables inside an n-dimensional rectangle. This algorithm is an extension of the univariate case which constructs a polynomial sequence, and measures sign variations of this sequence at the endpoints of the interval. We restrict ourselves to the case when $n = 2$.

Given two polynomials, $f_1(s,t)$ and $f_2(s,t)$, we construct the *volume function*, $V(u,s,t)$, as follows:

$$V(u,s,t) = \frac{Res_{a_2}(Res_{a_1}(f_1(a_1,a_2),f_3), Res_{a_1}(f_2(a_1,a_2),f_3))}{u^{deg(f_1(s,0))deg(f_2(s,0))}},$$

where $f_3(u,s,t,a_1,a_2) = u + (s-a_1)(t-a_2)$, Res_x refers to the resultant of two polynomials after eliminating x, and deg refers to the degree of the polynomial. We use the Sylvester resultant [Sal85] to eliminate one variable from two polynomials.

Given a square-free polynomial $p(x)$ we can construct a Sturm sequence of polynomials
$S_i = -remainder(S_{i-2}(x), S_{i-1}(x))$, where $S_1(x) = p(x)$ and $S_2(x) = p'(x)$. Treating the volume function V as a univariate polynomial in u, we construct its Sturm sequence $S_i(u,s,t)$. The Sturm sequence is specialized at $u = 0$ to give a sequence of bivariate polynomials $M(s,t)$.

Definition 2 *Given a sequence of polynomials $M(s,t)$ of length n, the $\mathbf{V}$ operator at (a_1,a_2) ($\mathbf{V}(M(a_1,a_2))$) gives the number of sign changes between consecutive terms of the sequence evaluated at (a_1,a_2). Correspondingly, the $\mathbf{P}$ operator is defined as $\mathbf{P}(M(a_1,a_2)) = n-1-\mathbf{V}(M(a_1,a_2))$.*

Given the bivariate sequence $M(s,t)$ and a rational axis aligned rectangle $\Gamma = [a_1,b_1] \times [a_2,b_2]$, the number of real roots of f_1 and f_2 inside Γ is given by

$$\frac{\mathbf{P}(M(b_1,b_2)) + \mathbf{P}(M(a_1,a_2)) - \mathbf{P}(M(b_1,a_2)) - \mathbf{P}(M(a_1,b_2))}{2}.$$

The justification for various steps and extension to arbitrary dimensions can be found in [Mil92].

2.4 Topological resolution of algebraic plane curves

The intersection curve between two surfaces is typically a high degree algebraic curve. In practice, it may have multiple real components. Topological resolution involves identifying critical points like turning points and singularities and establishing a unique connectivity between them. A number of efficient (poly-log time) algorithms have been developed for special kinds of algebraic curves. We use the algorithm by [AF88] for regular curves. The algorithm initially computes all the turning points of the curve. This is achieved in our case by taking partial derivatives and solving for common roots with the original curve inside a rational rectangle. A crucial step in establishing connectivity between the various turning points is to find the sign of the slope of the curve at certain irrational (algebraic) points. We present an exact algorithm to perform this step in section 3.3.2.

The identification of turning points divides the intersection curve into a set of monotonic curve segments. For each of these segments, we compute a bounding box (see Fig. 3) around it. Bounding boxes are needed to distinguish two such curve segments represented by the same algebraic equation. However, as seen from Fig. 3 , not all the bounding boxes are non-overlapping. We perform a subdivision of these boxes until these two criteria are satisfied. We use an algorithm described in [KM94] to perform this subdivision. The bounding boxes defined here are used to classify a point efficiently with respect to a curve.

3 Algorithm outline and Kernel routines

In this section, we give a brief overview of our algorithm and identify some steps that are susceptible to failure while using finite precision arithmetic. Finally, we describe the kernel routines and present efficient algorithms using exact arithmetic.

3.1 Computation of Boolean operation

The general outline for performing a Boolean operation between two solids follows. The overall approach is decomposed into two stages.

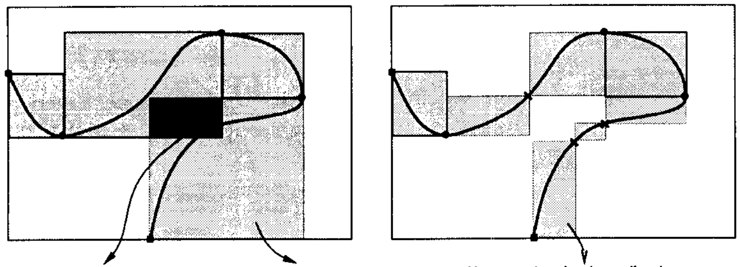

Figure 3: Bounding boxes around monotonic curve segments

I . Intersection curve computation (for each pair of patches):

1. Obtain the intersection curve(s) between the two untrimmed patches.
2. Find the points where the intersection curve meets the patch boundary.
3. Decompose the intersection curve into a set of monotonic curve segments.
4. Find the points where the intersection curve meets the trimming boundary, and subdivide the trimming and intersection curves, generating adjacency information.

II . Curve merging and boundary computation:

1. Merge intersection curves together in each patch.
2. Partition the patch domain.
3. Compute the adjacency graph and components separated by intersection curves.
4. Shoot rays from within one solid toward the other solid to classify components as inside or outside.
5. Propagate the information from step 3 in the adjacency graph to compute the boundary of the resulting solid.

3.2 Need for exact arithmetic

The use of exact arithmetic has been shown to be useful (and probably necessary) for dealing with degeneracies in the polyhedral domain ([For95]), and so it is likely to be needed in the non-linear domain. In addition, the use of finite precision arithmetic can result in numerical problems for non-degenerate cases.

We shall now identify two areas where our algorithm is susceptible to failure when using floating point arithmetic. In order to prevent these failures, we use exact arithmetic. Most of these errors finally boil down to either point orientation tests or comparison between two floating point numbers.

3.2.1 Computing trimmed patch intersections

Most algorithms using floating point arithmetic for computing the intersection curve use techniques like curve tracing or subdivision. As a result, these curves are approximated as piecewise linear curves or splines to within a fixed tolerance (which is either too conservative or arbitrarily chosen), or as algebraic curves with floating point coefficients. Since most of the surface patches we are dealing with are trimmed, we need to compute portions of the intersection curve that lie inside the trimmed boundaries of both the patches. Fig. 4(a) shows one such example. The curve $\mathbf{I}$ shown in dotted lines is the intersection curve in both the domains. $\mathbf{I}_0$ and $\mathbf{I}_1$ are the intersection curves on the left patch obtained from other surfaces. To compute the actual intersection curve for trimmed patches, we need to compute the intersection points of the curve with the trimming boundary. $\mathbf{p}_0$, $\mathbf{p}_1$, $\mathbf{p}_2$ and $\mathbf{p}_3$ are four such points on the right patch. If the boundary curves or the intersection curve are not accurate, neither are the $\mathbf{p}_i$'s. They may not even lie on the actual intersection curve. Corresponding to the $\mathbf{p}_i$'s, we need to compute $\mathbf{q}_i$'s on the other patch to determine which portions of the intersection curve to retain. This process is called *inversion*. Two problems can arise in inversion: (a) there may not be any corresponding point on the other patch (because $\mathbf{p}_i$'s do not lie exactly on the intersection curve), or (b) the $\mathbf{q}_i$'s could be positioned such that the curve segments $\mathbf{q}_0\mathbf{q}_1$ and $\mathbf{q}_2\mathbf{q}_3$ do not match up with $\mathbf{I}_0$ and $\mathbf{I}_1$ for curve merging. It is hard to perform this computation reliably using floating point arithmetic.

3.2.2 Component classification

Another area where floating point errors result in failure of the algorithm is during component classification. As we will describe later, we use ray shooting for this purpose and not sign evaluation of determinants as done by polyhedral modelers. The entire computation boils down to classifying whether a point lies inside or outside the trimming region. Fig. 4(b) shows an example. In most cases, classifying points like $\mathbf{q}_1$ is not a problem. One ray-shooting query will determine it. However, consider a point like $\mathbf{q}_0$

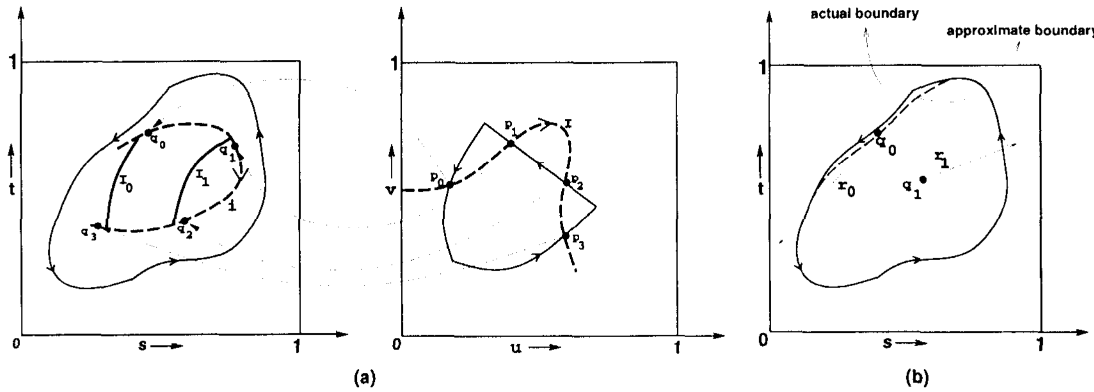

Figure 4: (a) Inaccurate point inversion for curve merging (b) Inaccurate point classification

which lies very close to the boundary. Approximate representations of the trimming boundary makes classifying q_0 a major problem. Depending on the choice of ray directions and the tolerances used we may get different classifications. This error could result in topologically inconsistent answers.

There are a number of similar problems which plague floating point modelers, and resolving these situations is no different from the ones highlighted. We believe that using exact arithmetic and representation is essential for reliable B-rep computation.

3.3 Exact computation of kernel routines

In this section, we will present efficient algorithms to implement the kernel routines in exact arithmetic. In particular, given algebraic curves and points, we present efficient algorithms for comparing two algebraic numbers, evaluating signs of slopes for resolution of regular algebraic curves, and classifying points with respect to a region.

3.3.1 Comparison between algebraic numbers

It was mentioned in Section 2.1 that each of the vertices in the solid are defined as roots of a set of polynomial equations with rational coefficients. Since we are dealing with rational parametric surfaces, each of these equations is bivariate. A vertex in the patch domain is the common solution of two equations, $f_1(s,t) = 0$ and $f_2(s,t) = 0$. This is usually an irrational number which cannot be represented exactly using floating-point arithmetic. In our algorithm, we represent each algebraic number as a small (based on lazy evaluation) rational rectangle. The rational rectangle is guaranteed to isolate each common root of $f_1(s,t)$ and $f_2(s,t)$ (taking into account the multiplicities of roots).

Consider a rational rectangle, Γ, that contains two algebraic numbers, γ_1 and γ_2 (see Fig. 5(a)) . Let γ_1 be the common root of $f_1(s,t)$ and $f_2(s,t)$, and γ_2 the common root of $g_1(s,t)$ and $g_2(s,t)$. There is an exact procedure to answer the question of equality of γ_1 and γ_2. Consider the new polynomial $p(s,t) = f_1^2(s,t) + f_2^2(s,t) + g_1^2(s,t) + g_2^2(s,t)$, and one of the original polynomials (say $f_1(s,t)$). It is easy to show that $p(s,t)$ and $f_1(s,t)$ have a common isolated root in Γ iff $\gamma_1 = \gamma_2$. We check for the common root using Sturm sequences. Further, this method can be extended to arbitrary dimensions.

3.3.2 Slope sign computation at algebraic points

Fig. 5(b) shows an example of an algebraic curve in a small region. Once the turning points are computed, they are separated using rational grid lines (s_1, s_2, t_1, and t_2 in the figure). In order to determine the topology of the curve (see Section 2.4), we need to compute the sign of the slope of the curve at points where the grid lines intersect the curve (e.g. p_4 in figure). Since p_4 is algebraic (a root of a univariate polynomial), we can compute it only within an interval (say, $[t_{41}, t_{42}]$). It is non-trivial to compute the sign of the slope directly (equivalent to computing signs of partial derivatives) at p_4. Let $f(s,t)$ be the equation of the curve. We know that there is a unique root of $f(s_1,t)$ in the interval $[t_{41}, t_{42}]$. Let $g(t) = f_s(s_1,t)$. We know that p_4 is not a common root of $f(s_1,t)$ and $g(t)$ (because it is not a turning point). We refine the interval $[t_{41}, t_{42}]$ such that there is no root of $g(t)$ within the interval. The sign of $g(t)$ can be obtained by evaluating it at any rational point in the interval.

3.3.3 Point classification

Classifying a component with respect to a solid amounts to classifying a point with respect to the trimmed domain. Problems associated with point classification using floating point arithmetic were highlighted earlier. We now describe our algorithm to exactly check whether a point (with algebraic number coordinates) lies inside or outside the trimming boundary. Initially, we assume that the actual point does not lie exactly on any algebraic curve that is part of the trimming boundary. This would be a case of four surfaces intersecting at a point and will be discussed when we handle degeneracies.

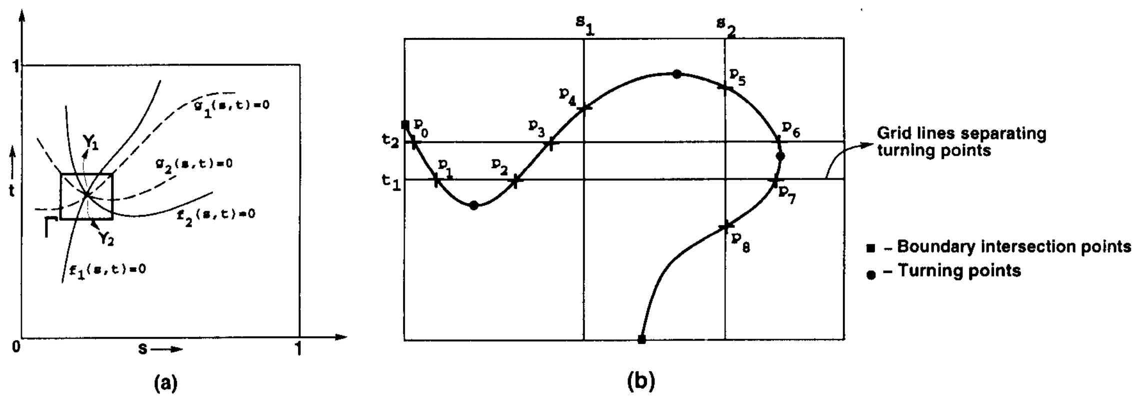

Figure 5: (a) Algebraic number comparison (b) Topological resolution of algebraic curves in finite domain

Given that the point is an algebraic number, we represent it with a rational rectangle of small size. We must ensure that the rectangle lies in at most one of the bounding boxes of the trimming curve (see Section 2.4). If it lies in more than one, the rectangle should be further refined so that it lies in only one. Each of the four vertices of the rectangle is classified with respect to the trimming boundary. This classification is done using a ray shooting technique (see Fig. 4(b)). To determine if a point (say, q_1 in figure) lies inside or outside the trimming boundary, we shoot an arbitrary semi-infinite ray (say, r_1 in figure) from it. The parity of the number of intersections of the ray with the boundary is sufficient to classify the point.

We perform this test on each of the four vertices of the rectangle. If all the results are same, actual point classification is done. The more interesting case is when some vertices of the rectangle yield different results. In this case, we refine our rectangle until consistency is achieved. We know this is assured because the point does not lie on any boundary curve. In practice, if we cannot classify the point after a few levels of refinement, we choose a different ray.

4 Exact B-rep Computation Algorithm

In this section, we briefly describe each step of our algorithm. It is implemented on top of the kernel routines.

4.1 Obtain intersection curve for two patches

We want to find the intersection curve in the domain of each of the two patches. To find the intersection curve in the domain of patch 1, we substitute the parameterization of patch 1 into the (precomputed) implicit representation of patch 2. A similar procedure obtains the intersection curve in the domain of patch 2.

4.2 Clip to Patch Boundary

Curve-surface intersection is used to find the points where the intersection curve meets the patch boundary. We want to find the intersection points of the curve and surface in both domains. Finding the intersection points of a curve with the boundaries of its own patch is straightforward - just substitute in the value for s or t and use a univariate Sturm sequence to isolate the roots. The *inversion* problem is to find the corresponding points in the domain of patch 2.

The inversion problem can be solved by considering the patch boundary as a curve in space. We compute the intersections of this space curve with the other patch (a curve-surface intersection). For example, consider the domain boundary $t = 0$. Then, the space curve corresponding to the patch boundary will be given by $X(s,0), Y(s,0), Z(s,0)$, and $W(s,0)$, where X, Y, Z, W give the parameterization of the first patch. The second patch will have the parameterization $\overline{X}(u,v), \overline{Y}(u,v), \overline{Z}(u,v), \overline{W}(u,v)$. Then, to solve for the points of intersection we need to solve for solutions of the equations (5) in their respective domains.

$$\begin{aligned} \overline{X}(u,v)\ W(s,0) &- X(s,0)\ \overline{W}(u,v) = 0 \\ \overline{Y}(u,v)\ W(s,0) &- Y(s,0)\ \overline{W}(u,v) = 0 \\ \overline{Z}(u,v)\ W(s,0) &- Z(s,0)\ \overline{W}(u,v) = 0 \end{aligned} \tag{4}$$

These equations may be solved using a trivariate Sturm sequence as described in [Mil92]. This will give solutions bounded in u, v, and s. However, the volume function computation involves three successive elimination steps (Sylvester resultant). Depending on the size of the coefficients, this process could be computationally intensive.

A more practical, though not exact, approach to solve this problem is to isolate all the roots in the s domain using a univariate Sturm sequence. This is followed by eliminating s from (5) to produce two independent equations in u and v. This is a bivariate Sturm sequence problem and is solved by a single Sylvester resultant computation. This gives all the solutions in (u, v) space. Determining the correspondence between the (u, v) pairs and s roots is done by comparing each of them in 3-space. We found that in practice, this method was significantly more efficient.

4.3 Decompose into monotonic curve segments

The process for isolating turning points and decomposing the intersection curve into segments monotonic in s and t (the patch parameters) was discussed in Section 2.4.

4.4 Prune intersection curves

We now have to trim the curve based upon the trimming boundary. Basically, we need to intersect the intersection curves with the trimming curves (represented as algebraic curves) and throw away sections of the intersection curve which are outside the trimming boundary.

Finding the points of intersection between the trimming curve and the intersection curve is relatively simple - use the bivariate Sturm sequence again. It is also relatively simple to find the corresponding points on the other patch - the trimming curve has a surface associated with it, and this surface intersected with the second patch gives another curve in the domain of that second patch, From this we obtain the intersection points on the intersection curve in the second patch, and figure out which points correspond by matching the intervals in 3-space. The actual pruning step is carried out by determining the orientation (inside/outside) of one point (we choose the starting point on the intersection curve) using 2D ray-shooting. Propagating this information to adjacent sections of the curve clearly identifies the curve segments that lie inside the trimmed region, and which lie outside. Along with this intersection curve segment, we also maintain the patch number of the other solid that defined this curve. We use this later when we update the topology information.

4.5 Merging intersection curves

In the previous step, we obtained the intersection curves on each patch for all patch-patch pairs. We now need to merge these curves (which will define the new trimming curves). This is done by matching the endpoints of each of the intersection curves, thereby partitioning the patch domain into closed loops. Notice that the monotonic segments of each intersection curve are already merged. Basically, we must check both endpoints of each intersection curve against the endpoints of all of the other intersection curves. If we find that two curves share a common endpoint, then we store this information and consider the curves merged. Checking for point equality is implemented as a kernel routine.

Once the intersection curves have been merged, we need to once again check whether the bounding boxes of the monotonic segments are overlapping and, if so, subdivide the curve appropriately.

4.6 Partitioning Trimming Boundaries

Once all the intersection curves are merged within each patch, they will partition the trimmed domain (if the assumptions that the individual solid boundaries are closed and compact are maintained). Otherwise it is a degenerate intersection (we discuss such cases in Section 6). Fig. 6(a) shows intersection curves inside a trimmed domain. $\mathbf{c_i}$'s (with endpoints $\mathbf{p_i}$ and $\mathbf{p_{i+1}}$) are monotonic curves (in both s and t) that form the trimmed boundary of the patch. $\mathbf{I_0}$, $\mathbf{I_1}$, and $\mathbf{I_2}$ are the intersection curves computed with various patches of the other solid. $\mathbf{t_0}$ is a turning point on the curve $\mathbf{I_2}$. As described earlier, all the turning points are identified before the topology of the algebraic intersection curve can be resolved. $\mathbf{q_i}$'s are points on the intersection curve where the curve intersects the trimmed boundary. Given this information, Fig. 6(b) shows the actual partitions ($\mathbf{R_i}$s). To compute the explicit B-rep of the resulting solid, each of these partitions is generated. We now present an algorithm that computes these partitions provided the intersection curves have no singularity in the trimmed domain.

The main idea in this algorithm is the fact that since the intersection curve segments (I_0 and I_1 in Fig. 6(c)) do not cross each other, each resulting partition starts at one endpoint of a curve segment, and ends at the other endpoint of the same curve segment. We shall assume that the trimming curves and the intersection curves are given in a specific order. We number the endpoints of the intersection curve segments such that $\mathbf{q_{2j}}$ and $\mathbf{q_{2j+1}}$ belong to $\mathbf{I_j}$. The algorithm works in three steps.

- Each endpoint of a curve segment (for example, $\mathbf{q_0}$ of $\mathbf{I_0}$) lies on a unique curve (except when it coincides with one of the curve endpoints of the boundary) of the trimming boundary. In fact, points like $\mathbf{q_0}$ are determined as the intersection of $\mathbf{I_0}$ with $\mathbf{c_0}$. Note that even though $\mathbf{c_0}$ and $\mathbf{c_1}$ could be part of the same algebraic curve, the association of $\mathbf{q_0}$ with $\mathbf{c_0}$ is determined because of the monotonicity of the $\mathbf{c_i}$'s. Each boundary curve $\mathbf{c_i}$ is then partitioned into multiple segments depending on the number $\mathbf{q_j}$'s lying on it.
- This is followed by a traversal of the trimming boundary in a consistent order by maintaining a stack. Two types of elements are pushed in the stack - curve segments, and curve endpoints. Initially, we keep pushing in the boundary curve segments until we reach a vertex like $\mathbf{q_0}$. Let the vertex number be k. If the topmost curve endpoint type of the stack (say, l) has a number $(k+1)$ or $(k-1)$, then a partition has to be read out. Otherwise, vertex k is pushed into the stack followed by all the curves that comprise $I_{\lfloor k/2 \rfloor}$. If a decision to read out a region has been reached, all the curve segments until vertex l are popped. Curves comprising $I_{\lfloor k/2 \rfloor}$ are pushed again because they are required by the next region too. The order in which these curve segments are pushed into the stack has to be monitored carefully so that a region which is read out is oriented consistently.
- Till now, we have considered only intersection curve segments whose endpoints lie on the trimming boundary. However, there may be loops that lie completely

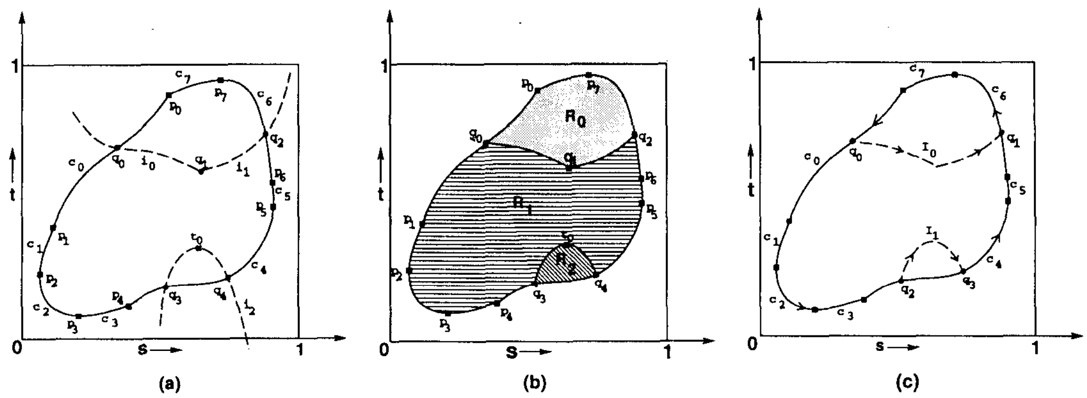

Figure 6: (a) Intersection curves inside trimmed domain (b) Partitions introduced by intersection curves (c) Partitioning a trimmed patch with chains of algebraic curves

inside the boundary. Any loop is present (if at all) inside one of the obtained partitions. Each of the loops (starting from the innermost if the loops are nested) themselves form a partition. The remaining part of the region (it has boundaries with multiple components) is broken into simple regions by introducing a simple cut from the loop to the boundary of the partition or the next loop.

This completes the algorithm to compute the partitions introduced by intersection curves. A feature of this algorithm is that the adjacency structure between the various partitions (which is necessary to avoid redundant, expensive ray-shooting queries during component classification) are obtained by the order in which they are read out.

4.7 Updating Topological Information

It is clear from the previous section that intersection computation introduces new vertices, edges, and faces in the solid. This change needs to be incorporated in our topological structure. Further, information about the adjacency between the various faces significantly reduces the component classification time. At this time, we just concentrate on the face adjacency. Vertex and edge adjacency are updated during final solid generation.

The new graph is a refinement of the original adjacency graph. Remember that a vertex of the graph corresponds to a face of the solid. Each vertex in the original graph is split into a few vertices depending on the partitions obtained due to the intersection curves. We need to figure out the adjacency relationship between the newly created vertices. Consider, for example, that vertices $\mathbf{u}$ and $\mathbf{v}$ were adjacent in the original graph. Due to the intersection curves, let the vertex $\mathbf{u}$ be split into $\mathbf{u}_1, \mathbf{u}_2, \ldots, \mathbf{u}_m$, and let the vertex $\mathbf{v}$ be split into $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n$. The adjacency between the various $\mathbf{u}_i$'s (similarly $\mathbf{v}_i$'s) has already been determined (during partitioning). These adjacencies (let us call it set S) are purposely left out in the new graph. Let e be the edge along which $\mathbf{u}$ and $\mathbf{v}$ were adjacent in the original graph, and let it be divided into k portions during partitioning. Then all the adjacencies between $\mathbf{u}_i$'s and $\mathbf{v}_j$'s can be obtained in $O(k)$ time. The number of connected components in this graph gives the number of solid components introduced by the intersection curves. Let the solid components be named $CC_0, CC_1, \ldots$. Note that each CC_i has a collection of faces.

To obtain the connectivity between the various CC_is, we introduce some notation. Let R be a mapping which takes a vertex in the new graph to the corresponding vertex in the original graph. For example, if $\mathbf{u}$ was split into $\mathbf{u}_1$, $\mathbf{u}_2$, $\ldots$, $\mathbf{u}_m$, then $R(\mathbf{u}_i) = \mathbf{u}$. Two components CC_i and CC_j are connected if

$$\{\exists \mathbf{u}_1 \in CC_i, \exists \mathbf{u}_2 \in CC_j | R(\mathbf{u}_1) = R(\mathbf{u}_2) \textit{ and } (\mathbf{u}_1, \mathbf{u}_2) \in S\}$$

Using this, we have obtained the various components and their connectivity. Next we resolve each of these components (inside/outside) with respect to the other solid.

4.8 Component Classification

Component classification involves determining whether a given component of one solid is inside or outside the other solid. It is obvious that the entire component (as obtained in the previous section) lies completely inside or outside the other solid. In most polyhedral modelers, component classification is carried out locally [Hof89]. When dealing with sculptured surfaces, though, the same technique cannot be used. The most general method used instead is based on ray-shooting. Ray-shooting is done by firing a semi-infinite ray in an arbitrary direction and checking for intersections with the other solid. If the number of intersections is even, the point (and hence, the entire component) lies outside the solid; if it is odd, it lies inside.

There are three steps involved in our algorithm to perform component classification. The first step involves getting a point that is part of the component. This reduces to finding a point inside the trimming boundary of one patch. This is accomplished by 2D ray-shooting. We initially choose some rational point $p = (s, t)$ in the domain such that t lies between the lower and upper extents of the trimming boundary. A horizontal line passing through p (in both directions) is intersected with the boundary, and all the intersections are determined using the root isolation method of Milne [Mil92]. The intersections must be even in number and are of the form $(s_1, t), (s_2, t), \ldots, (s_{2n}, t)$. The s_i's are algebraic numbers and are represented as small rational intervals. Choosing the midpoint of s_{2i-1} and s_{2i} for $i = 1, 2, \ldots, n$ gives one rational point inside the trimming boundary. Let this point be called **q**.

The second step involves actual ray-shooting in 3-space. To perform 3D ray-shooting, we compute q's mapping in 3-space using the patch parametrization. Let this point be (x_q, y_q, z_q). This point is rational because of the formulation of the patch parametrization. We pick a random direction and fire a semi-infinite ray in that direction. We compute all the intersections of each patch of the other solid with this ray. This is done similar to the curve-surface intersection computation described earlier in this section. However, not all the intersection points computed this way lie inside the trimmed boundary of the patch. Checking if the intersection point lies inside the trimmed boundary, as in Section 3.3.3, is the third step of our algorithm.

4.9 Final B-rep Generation

The trimmed patches that make up the final solid are determined by the Boolean operation performed. Given two solids $solid_1$ and $solid_2$, we decide on the final B-rep depending on the Boolean operation as follows:

- *Union:* All components of $solid_1$ that lie **outside** $solid_2$, and vice-versa are retained.
- *Intersection:* All components of $solid_1$ that lie **inside** $solid_2$, and vice-versa are retained.
- *Difference:* All components of $solid_1$ that lie **outside** $solid_2$, and all components of $solid_2$ that lie **inside** $solid_1$ are retained.

We also update the topology information. Each connected component that is retained in the final solid has some graph vertices (faces of the solid) whose complete adjacency is not determined. These correspond to edges which are formed by intersection curves. The vertex that should be adjacent to this vertex along this edge in the final graph is part of the other solid, and is the surface that formed the intersection curve. We maintain the patch number of the other solid that resulted in an intersection curve, and use it to complete the adjacency. From this graph, the entire topological information is easily computable.

5 Analysis and Performance

In this section, we shall discuss some of the theoretical and empirical complexity analysis for some important steps of our algorithm. It is clear that the most dominating steps in terms of time are the root isolation of bivariate polynomials, and the topological resolution of intersection curves. We optimize these algorithms, and implement them as part of the kernel routines. The complexity of steps involving partitioning based on the intersection curve, and the face connectivity generation are very small compared to the total cost, and their analysis is omitted here. We now give the theoretical worst-case complexity of the root isolation algorithm.

5.1 Worst case analysis

Root isolation: Most of the results involving real root isolation are based on Sturm sequences, and we quote a result from Davenport [Dav85] for the worst-case time complexity of root isolation algorithm for univariate polynomials.

Theorem 1 *[Dav85] The running time of the root isolation algorithm based on Sturm sequences of univariate polynomials is bounded by $O(n^6(\log n + \log \Sigma a_i^2)^3)$, where n is the degree of the polynomial and a_i are its coefficients.*

But this bound is too pessimistic, and a result based on [Hei71] predicts the average case to be more like $O(n^4)$.

Sylvester resultant: Given two polynomials of degree m and n respectively, the coefficient size of the resultant is bounded by $[(m + n)\ c]^{m+n}$ [KKM96], where c is the maximum coefficient size of the polynomials. For the case of quadrics, tori, and other low-degree solids that we deal with m and n is at most four. Thus the bit complexity of the coefficients of the univariate polynomial is roughly 8 times the original bit complexity in the worst case.

Topological resolution of intersection curves: The algorithm described [AF88] computes the sign invariant decomposition of an algebraic curve in a finite domain. The time complexity is given by the following result:

Theorem 2 *[AF88] Given a bivariate polynomial of total degree n and coefficient size d in E^2, it is possible to obtain a sign invariant decomposition of the curve in time $O(n^{12}\ (d + \log n)^2\ \log n)$.*

5.2 Improving performance of exact arithmetic

As exhibited by the worst case bounds, the arithmetic of these symbolic coefficients is expensive. Actually, the cost of arithmetic operations is quadratic in the size of its coefficients. Even though, in reality, we do not experience such a drastic increase in bit complexity for intermediate computation, it nevertheless grows. To reduce the cost of arithmetic operations, we perform all our computations over finite fields, and use a probabilistic algorithm based on the Chinese Remainder Theorem to recover the actual coefficients [MC93]. The time complexity of the resultant computation (using interpolation algorithm in [MC93]) is directly proportional to the number of primes used in the finite field computation. To reduce this number, we use primes of maximum possible magnitude. Most current implementations of ***bignum*** libraries use finite fields of order 2^{16} to prevent overflow when taking products. Most of the current machines provide multiplication instructions that give the result out in two registers. Taking advantage of this fact, we use an assembly level subroutine that performs multiplication. This allows us to use finite fields of order as high as 2^{31}. Compared to finite fields of order 2^{16}, we get almost two-fold improvement in speed.

Lazy evaluation during root isolation: Another optimization that we perform to improve our speedup is to evaluate rational intervals (in root isolation) as lazily as possible. This is based on the assumption that the worst case bounds that govern the closeness of roots actually occur very rarely in practice. Thus, most of the time, we are able to isolate the roots of polynomials inside the domain of surfaces quickly. However, during later computation of other roots, it might be necessary to make sure that two roots are not the same (because of large overlapping intervals). At this time, we refine the computed intervals further and try to isolate their intervals. This lazy approach to root isolation behaves almost like an output sensitive algorithm.

5.3 Empirical Performance

It is quite clear that the running time of our algorithm grows quickly as a function of the degree of the solids. We have tried computing B-reps for low degree solids like quadrics and tori. Based on our experience for these solids, *the use of exact arithmetic slows us down by slightly more than one order of magnitude* (as compared to B-rep algorithms based on IEEE floating point arithmetic). The performance varies with the degrees, relative orientation of two solids and the complexity of the output. In cases when the solids are close to being degenerate, the coefficient sizes grow more quickly, and the algorithm slows down. However, we feel that this penalty is justified for improving the accuracy of the modeler and performing reliable computation. We would like to mention that though we have implemented various steps of the algorithm independently, they have not yet been integrated together. Furthermore, our implementations have not been optimized (especially the kernel routines).

For most quadrics, the parameterization is biquadric (leading term s^2t^2). Substituting this into the implicit form of the other patch results in an intersection curve which is biquartic (s^4t^4) in the domain of the patch. While computing the resultants of two such equations for root isolation, the order of the resultant matrix is at most 8×8. For these orders of matrices, the resultant algorithm finishes in about 50 milliseconds. The algorithm uses the special structure of the Sylvester matrix, and evaluates the resultant using interpolation methods. The volume function (see Section 8.2), thus computed, is a univariate polynomial in u of maximum degree 8. Depending on the size of the actual coefficients, the time taken to compute the Sturm sequence ranges from 0.2 to 1 second. The entire root isolation procedure for these cases typically takes less than a second to 2 seconds. We estimate that for a boolean operation between two low-degree solids, our algorithm will compute the exact B-rep in $40-100$ seconds - slightly more than one order of magnitude slower than a finite precision implementation. Timings are based on an HP 712/100 model workstation.

5.4 Example

In this section, we illustrate some key parts of our approach using an example of two cylinders which are just interpenetrating (see Fig. 7(a)). The cylinders are of radius 1 and 0.5, and their centers are spaced 1.49 units apart. The cylinders are rotated with respect to each other. We divide the surface of each cylinder into four equal parts and represent each of them as a rational parametric surface with rational coefficients. The parametric form of a sample patch from each cylinder are given below.

$$\begin{aligned} X(s,t) &= 1 - s^2 \\ Y(s,t) &= 2s \\ Z(s,t) &= 2t + 2s^2t - 1 \\ W(s,t) &= 1 + s^2 \end{aligned}$$

$$\begin{aligned} \overline{X}(u,v) &= \tfrac{199}{100}u^2 + \tfrac{99}{100} \\ \overline{Y}(u,v) &= \tfrac{4}{5}u - \tfrac{6}{5}v - \tfrac{6}{5}u^2v + \tfrac{3}{5} \\ \overline{Z}(u,v) &= \tfrac{3}{5}u + \tfrac{8}{5}v + \tfrac{8}{5}u^2v - \tfrac{4}{5} \\ \overline{W}(u,v) &= 1 + u^2 \end{aligned}$$

After implicitizing using Dixon's formulation, the implicit forms are

$$\begin{aligned} f(x,y,z,w) &= w^2 - x^2 - y^2 \\ g(x,y,z,w) &= 19701\,w^2 - 29800\,x\,w + 10000\,x^2 + \\ &\quad 6400\,y^2 + 9600\,y\,z + 3600\,z^2 \end{aligned}$$

To obtain the intersection curve of the two patches in the domain of the first patch, we substitute its parameterization into the implicit form of the second patch ($g(x,y,z,w)=0$).

$$\begin{aligned} h(s,t) = &-3501 + 19200\,s - 45002\,s^2 - 59501\,s^4 + \\ &14400\,t - 38400\,st + 14400\,s^2t - \\ &38400\,s^3t - 14400\,t^2 - 28800\,s^2t^2 - \\ &14400\,s^4t^2 = 0 \end{aligned}$$

Since the patches are untrimmed, we have to compute the starting points of the curve on the boundary of the patch. Substituting $s = 0$ into h and computing the volume function for this univariate case, we get

$$\begin{aligned} V(u,t) = &-3501 + 14400\,t - 14400\,t^2 + \\ &14400\,u - 28800\,t\,u - 14400\,u^2 \end{aligned}$$

We computed the Sturm sequence of this volume function, and isolated the roots of the original equation between $[0,1]$ to within a precision of $\frac{1}{100}$. The two roots were

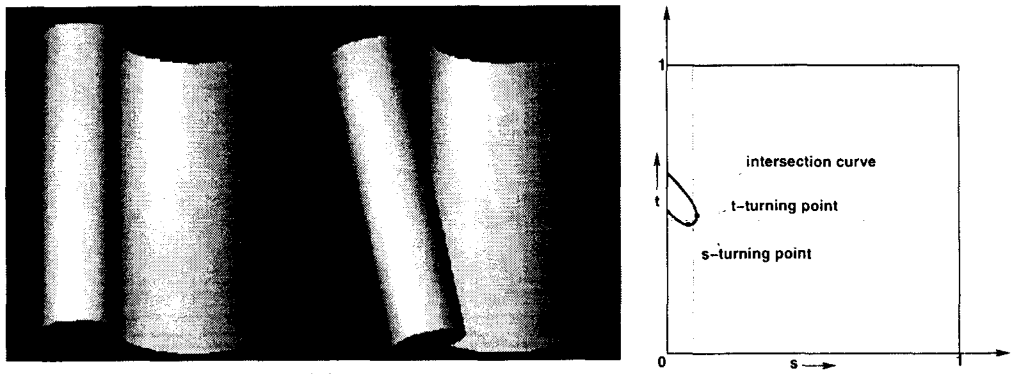

(a) (b)

Figure 7: (a) Two views of just interpenetrating cylinders (b) Complete intersection curve in the domain of one patch

$$\left[\frac{226834}{390625}, \frac{1838}{3125}\right], \quad \left[\frac{1298}{3125}, \frac{6634}{15625}\right]$$

These numbers give the ranges in t for which there is an intersection of the intersection curve with the s=0 edge of the patch domain. Of these, only the first one corresponds to a point inside the domain of the second patch. This was obtained by inversion and a bivariate Sturm sequence generation. The point corresponding to $\left(0, \left[\frac{226834}{390625}, \frac{1838}{3125}\right]\right)$ inside the domain of the second patch is $\left(\left[\frac{7352}{78125}, \frac{64}{625}\right], \left[\frac{43682}{78125}, \frac{8866}{15625}\right]\right)$. The s (or t) turning points on the intersection curve were obtained by performing bivariate Sturm sequence root isolation on the pairs of polynomials $h(s,t)$ and $h_s(s,t)$ ($h_t(s,t)$). The s and t turning points were found to be $\left(\left[\frac{20464}{390625}, \frac{4352}{78125}\right], \left[\frac{5764}{15625}, \frac{146044}{390625}\right]\right)$ and $\left(\left[\frac{1096}{15625}, \frac{248}{3125}\right], \left[\frac{2}{5}, \frac{6346}{15625}\right]\right)$ respectively.

Now that we have the turning points of the intersection curve, we compute the topological resolution. After separating the turning points and evaluating the signs of slope at various grid lines, the connectivity of the curve is obtained unambiguously. The resulting curve in the domain of the first patch is shown in Fig. 7(b). The intersections of the curve with the s=0 axis, the turning points in s and t, and the grid lines for curve connectivity computation are all shown in the picture. Note that the curve shown was obtained by intersecting this patch with *all* the patches of the second cylinder.

6 Degeneracies

A number of degenerate cases can arise when dealing with curved surfaces. Some of these degeneracies are of the same general type as is found in a polyhedral modeler, while some others arise only with curved surface modelers. These include:

- **Two surfaces meeting at a point:** This is a singularity which we assume does not occur. If it does occur, we can find it by noticing that the intersection curve has an s turning point and a t turning point at the same position.
- **Two surfaces meeting at a curve:** This is a degenerate case when the surfaces are tangent to each other along that curve. We will be able to detect this when we generate the adjacency graph by finding whether two components which should be adjacent are actually part of the same component.
- **Two surfaces overlapping:** This corresponds to a face-face overlap in the polyhedral domain. Here, though, if the surfaces we use have an irreducible implicit form, then they will not overlap unless this form is identical.
- **A surface just touching an edge:** This is an edge-face contact in the polyhedral domain, and can happen when three surfaces meet in a curve. In our representation, this will appear as an intersection curve which is tangent to a trimming curve (see Fig. 8(a)). Such a case can be automatically eliminated if we check *each* component of the intersection curve to see whether it is in the trimmed region. This does not allow us to use the speed-up of propagating the information about one component of the intersection curve to all other components of that curve.
- **Four surfaces meeting at a point:** This, is the foundation for several types of degeneracies and will be discussed next.

Examples of four surfaces meeting at a point include when a vertex of one solid lies on the surface of another solid, or when the edges of two solids meet. Obviously, the vertex can be thought of as the intersection of three surfaces, and the edges can be thought of as the intersection of two surfaces, thus the cases mentioned would involve the intersection of four surfaces.

Even more degenerate cases, such as two vertices meeting, or a vertex lying on an edge, are possible, but these

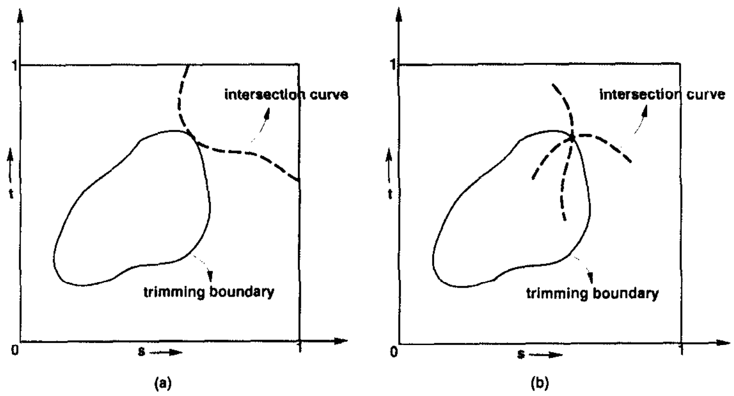

Figure 8: (a) Surface-edge contact degeneracy (b) Four surfaces meeting at a point

can be viewed as 5 or 6 surfaces meeting at a point - i.e. at least four surfaces are still meeting at a point.

These cases will manifest themselves in our modeler as three (or more) curves meeting at a common point in the domain of some patch (see Fig. 8(b)). Assume these three curves are $f1$, $f2$, and $f3$. We can find out whether this case has occurred by checking equality of the intersection of $f1$ and $f2$ in some interval with the intersection of $f1$ and $f3$ (or $f2$ and $f3$) in that same interval.

Degeneracies in the polyhedral case can generally be classified into the category of four planes meeting at a point. It has been shown [For95] that a simple perturbation scheme applied to a single basic geometric predicate can eliminate these degeneracies. No obvious extension of this method exists in the curved surface domain, though there is hope that some perturbation method can be developed which would work similarly.

7 Extensions and Future Work

In this paper, we have presented representations and algorithms to compute B-reps for boolean combinations of low-degree solids specified with rational parametric surfaces. We use exact arithmetic to perform reliable computations on a number of kernel routines, upon which the rest of the modeler is built. The efficient and accurate implementation of the kernel routines allows us to have an efficient and reliable method for the overall B-rep computation.

There are a number of ways in which the method we have described might be extended. These are:

- We need to be able to handle singularities in intersection curves and surfaces.
- In order to claim robustness, we will need to be able to deal with all degneracies (possibly using perturbations).
- Although this method is efficient for lower-degree surfaces, it would be useful for it to be more efficient for higher degree surfaces.
- It would also be useful to extend our method to deal with non-manifold cases or with non-parametric/non-algebraic surfaces.
- Finally, we would like to investigate the use of parallelism or a combination of floating-point and exact arithmetic to get a faster implementation.

References

[AF88] S. Arnborg and H. Feng. Algebraic decomposition of regular curves. *Journal of Symbolic Computation*, 5:131–140, 1988.

[BMP94] M. Benouamer, D. Michelucci, and B. Peroche. Error-free boundary evaluation based on a lazy rational arithmetic: a detailed implementation. *Computer-Aided Design*, 26(6), 1994.

[Bra75] I. Braid. The synthesis of solid bounded by many faces. *Comm. ACM*, 18:209-216, 1975.

[Can88] J.F. Canny. *The Complexity of Robot Motion Planning.* ACM Doctoral Dissertation Award. MIT Press, 1988.

[CB89] M. S. Casale and J. E. Bobrow. A set operation algorithm for sculptured solids modeled with trimmed patches. *Computer Aided Geometric Design*, 6:235–247, 1989.

[Dav85] J. H. Davenport. Computer algebra for cylindrical algebraic decomposition. Technical Report TRITA-NA-8511, NADA, KTH, Stockholm, 1985.

[Dix08] A.L. Dixon. The eliminant of three quantics in two independent variables. *Proceedings of London Mathematical Society*, 6:49–69, 209–236, 1908.

[FBZ93] S. Fang, B. Bruderlin, and X. Zhu. Robustness in solid modeling: a tolerance-based intuitionistic approach. *Computer-Aided Design*, 25(9):567–576, 1993.

[For95] S. Fortune. Polyhedral modeling with exact arithmetic. *Proceedings of ACM Solid Modeling*, pages 225–234, 1995.

[For96] S. Fortune. Robustness issues in geometric algorithms. In M.C. Lin and D. Manocha, editors, *Applied Computational Geometry*, pages 9–14. Springer-Verlag, 1996.

[Hea95] M. Higashi and et al. Face-based data structure and its application to robust geometric modeling. *Proceedings of ACM Solid Modeling*, pages 235–246, 1995.

[Hei71] L. E. Heindel. Integer arithmetic algorithm for polynomial real zero determination. *Journal of ACM*, 18(4):535–548, 1971.

[HHK89] C. Hoffmann, J. Hopcroft, and M. Karasick. Robust set operations on polyhedral solids. *IEEE Computer Graphics and Applications*, 9(6):50–59, 1989.

[Hof89] C.M. Hoffmann. *Geometric and Solid Modeling*. Morgan Kaufmann, San Mateo, California, 1989.

[Hof96] C.M. Hoffmann. How solid is solid modeling. In M.C. Lin and D. Manocha, editors, *Applied Computational Geometry*, pages 1–8. Springer-Verlag, 1996.

[Jac95] D. Jackson. Boundary representation modeling with local tolerances. *Proceedings of ACM Solid Modeling*, pages 247–253, 1995.

[KKM96] J. Keyser, S. Krishnan, and D. Manocha. Efficient boundary generation of low-degree sculptured solids using exact arithmetic. Technical Report TR96-040, Department of Computer Science, University of North Carolina, 1996.

[KM94] S. Krishnan and D. Manocha. An efficient surface intersection algorithm based on the lower dimensional formulation. Technical Report TR94-062, Department of Computer Science, University of North Carolina, 1994. To appear in ACM Trans. on Computer Graphics.

[KM96] S. Krishnan and D. Manocha. Efficient representations and techniques for computing b-rep's of csg models with nurbs primitives. In *Proceedings of CSG'96*, pages 101–122. Information Geometers Ltd, 1996.

[Mac02] F.S. Macaulay. On some formula in elimination. *Proceedings of London Mathematical Society*, 1(33):3–27, May 1902.

[Man88] M. Mantyla. *An Introduction to Solid Modeling*. Computer Science Press, Rockville, Maryland, 1988.

[MC93] D. Manocha and J.F. Canny. Multipolynomial resultant algorithms. *Journal of Symbolic Computation*, 15(2):99–122, 1993.

[MG91] J. Miller and R. Goldman. Combining algebraic rigor with geometric robustness for the detection and calculation of conic sections in the intersection of two quadric surfaces. *Proceedings of ACM Solid Modeling*, pages 221–233, 1991.

[Mil92] P. S. Milne. On the solutions of a set of polynomial equations. In *Symbolic and Numerical Computation for Artificial Intelligence*, pages 89–102, 1992.

[RV85] A.A.G. Requicha and H.B. Voelcker. Boolean operations in solid modeling: boundary evaluation and merging algorithms. *Proceedings of the IEEE*, 73(1), 1985.

[Sal85] G. Salmon. *Lessons Introductory to the Modern Higher Algebra*. G.E. Stechert & Co., New York, 1885.

[Sar83] R F Sarraga. Algebraic methods for intersection. *Computer Vision, Graphics and Image Processing*, 22:222–238, 1983.

[Sed83] T.W. Sederberg. *Implicit and Parametric Curves and Surfaces*. PhD thesis, Purdue University, 1983.

[Seg90] M. Segal. Using tolerances to guarantee valid polyhedral modeling results. In *Proceedings of ACM Siggraph*, pages 105–114, 1990.

[SI89] K. Sugihara and M. Iri. A solid modeling system free from topological inconsistencis. *J. Inf. Proc., Inf. Proc. Soc. of Japan*, 12(4):380–393, 1989.

[Yap90] C. Yap. A geometric consistency theorem for a symbolic perturbation scheme. *Journal of Symbolic Computation*, 40:2–18, 1990.

[Yu92] J. Yu. *Exact arithmetic solid modeling*. PhD thesis, Purdue University, 1992.