

Transfinite interpolation for well-definition in error analysis in solid modelling

N. F. Stewart* and M. Zidani

Département IRO, Université de Montréal,
C.P. 6128, Succ. CentreVille, Montréal, Qc, H3C 3J7, Canada.
{stewart, zidanima}@iro.umontreal.ca

Abstract. An overall approach to the problem of error analysis in the context of solid modelling, analogous to the standard forward/backward error analysis of Numerical Analysis, was described in a recent paper by Hoffmann and Stewart. An important subproblem within this overall approach is the well-definition of the sets specified by inconsistent data. These inconsistencies may come from the use of finite-precision real-number arithmetic, from the use of low-degree curves to approximate boundaries, or from terminating an infinite convergent (subdivision) process after only a finite number of steps.

An earlier paper, by Andersson and the present authors, showed how to resolve this problem of well-definition, in the context of standard trimmed-NURBS representations, by using the Whitney Extension Theorem. In this paper we will show how an analogous approach can be used in the context of trimmed surfaces based on combined-subdivision representations, such as those proposed by Litke, Levin and Schröder.

A further component of the problem of well-definition is ensuring that adjacent patches in a representation do not have extraneous intersections. (Here, ‘extraneous intersections’ refers to intersections, between two patches forming part of the boundary, other than prescribed intersections along a common edge or at a common vertex.) The paper also describes the derivation of a bound for normal vectors that can be used for this purpose. This bound is relevant both in the case of trimmed-NURBS representations, and in the case of combined subdivision with trimming.

1 Introduction

One of the fundamental problems in proving rigorous theorems in the area of robustness of numerical methods, in the field of solid modelling, is that the data normally provided to the algorithm is not only in error, it may be fundamentally inconsistent. These inconsistencies, in the data purportedly defining a set to be manipulated by a solid-modelling algorithm, come from the use of finite-precision real-number arithmetic, from the use of low-degree curves to approximate boundaries, or from terminating an infinite convergent (subdivision)

* The research of the first author was supported in part by a grant from the Natural Sciences and Engineering Research Council of Canada.

process after only a finite number of steps. Representations often have both a topological component and a geometric component; the geometric component may itself be internally inconsistent, and it may be inconsistent with the topological component. One approach to resolving these inconsistencies is to propose a definition of a set, satisfying strong guarantees of proximity to the given inconsistent input data, and to take this as specifying the input set. This approach permits subsequent rigorous proof, of theorems concerning algorithms that manipulate sets, in terms of well-defined input sets.

It should be noted that we are concerned with the case when the input data is uncertain and, as mentioned above, possibly inconsistent.

In [1] it was shown how the Whitney Extension Theorem [2] can be used to perform transfinite interpolation in order to realize the above goal of well-definition of a given input set, in the context of standard trimmed-NURBS representations. In this paper we will show how it can be used in an analogous way for the well-definition of sets defined by combined subdivision surfaces [3]. Then, later in the paper, we will consider a more general setting, including both of the special cases just mentioned, and provide a result for bounding normal vectors. This result can be used to ensure that adjacent patches in a representation do not have extraneous intersections. (Here, ‘extraneous intersections’ refers to intersections, between two patches forming part of the boundary, other than prescribed intersections along a common edge or at a common vertex.)

Assuring the well-definition of input sets is an important subproblem within the forward/backward error analysis described by Hoffmann and Stewart [4].

2 Transfinite interpolation

A transfinite interpolant is a surface that matches data on the entire boundary of a two-dimensional domain, rather than just at a finite number of points. Such surfaces can be obtained, for example, as solutions of the Dirichlet problem, minimizing the functional

$$\iint_D (f_x^2 + f_y^2) dx dy \quad / \quad \iint_D \|f\|^2 dx dy$$

under the Dirichlet boundary conditions [5, p. 110], or by finding area-minimizing solutions. The shape of such solutions is illustrated by a soap-film stretched over a wire-frame (see for example [6, frontispiece]).

Transfinite interpolation has been used since the earliest days of geometric and solid modelling. For example, the Coons patch [7] is a C^1 -continuous transfinite interpolant; see also [8]. More recently this kind of interpolation has been used by Gross and Farin to generalize Sibson’s interpolant to the case of boundary interpolation [9], and by Shapiro and his students [10] in the context of modelling heterogeneous materials on a point-by-point basis. As mentioned in the Introduction, it has also been used in quite a different way, to provide a definition of a well-formed set in the study of robustness [1, 4]. In this last-mentioned context, the transfinite interpolant is not actually computed, but is

introduced only to permit proof of rigorous theorems about a single well-defined set that can be viewed as the one specified by the inconsistent data provided to a numerical method. This approach will be described in the following section.

3 Whitney extension to provide transfinite interpolation in the context of solid modelling

We begin this section by stating a version [1] of the Whitney Extension Theorem [2] that is appropriate for our purposes.

Suppose that we have a mapping

$$\epsilon : C \rightarrow R^3, \quad (1)$$

where C is a compact subset of R^2 , and suppose that each component ϵ of ϵ satisfies a Lipschitz condition

$$|\epsilon(\mathbf{p}_1) - \epsilon(\mathbf{p}_2)| \leq L \cdot \|\mathbf{p}_1 - \mathbf{p}_2\|, \quad \mathbf{p}_1, \mathbf{p}_2 \in C$$

where

$$L = \sup_{\mathbf{p}_1, \mathbf{p}_2 \in C, \mathbf{p}_1 \neq \mathbf{p}_2} \frac{|\epsilon(\mathbf{p}_1) - \epsilon(\mathbf{p}_2)|}{\|\mathbf{p}_1 - \mathbf{p}_2\|} \quad (2)$$

is finite. Then [2] we can extend ϵ to a continuous function on all of R^2 that satisfies a Lipschitz condition with the same Lipschitz constant L .

In fact, let

$$l(\mathbf{p}) = \sup_{\mathbf{q} \in C} \{\epsilon(\mathbf{q}) - L \cdot \|\mathbf{p} - \mathbf{q}\|\}, \quad \mathbf{p} \in R^2, \quad (3)$$

$$u(\mathbf{p}) = \inf_{\mathbf{q} \in C} \{\epsilon(\mathbf{q}) + L \cdot \|\mathbf{p} - \mathbf{q}\|\}, \quad \mathbf{p} \in R^2, \quad (4)$$

and

$$\epsilon(\mathbf{p}) = \frac{1}{2}[l(\mathbf{p}) + u(\mathbf{p})], \quad \mathbf{p} \in R^2. \quad (5)$$

It is easy to show [11] that any continuous function satisfying the Lipschitz condition on R^2 is bracketed by the *lower function*, $l(\mathbf{p})$, and the *upper function*, $u(\mathbf{p})$, and [12] that $l(\mathbf{p})$ and $u(\mathbf{p})$ are themselves solutions to the extension problem. Furthermore [1], $\epsilon(\mathbf{p})$ is a solution to the extension problem satisfying

$$|\epsilon(\mathbf{p})| \leq \sup_{\mathbf{q} \in C} |\epsilon(\mathbf{q})|, \quad \mathbf{p} \in R^2.$$

Theorem 1. (Whitney, 1934) *The mapping $\epsilon(\mathbf{p})$ given in (5) is a continuous function on R^2 that coincides with the mapping (1) given on C . Furthermore, the mapping given in (5) satisfies a Lipschitz condition everywhere in R^2 , with Lipschitz constant L defined as in (2), and it satisfies the inequality*

$$|\epsilon(\mathbf{p})| \leq \sup_{\mathbf{q} \in C} |\epsilon(\mathbf{q})|, \quad \mathbf{p} \in R^2.$$

The Lipschitz constants for each of the components of ϵ can be used to define a Lipschitz constant for the vector-valued function ϵ .

3.1 The trimmed-NURBS case

In this subsection we give a summary of the use of Whitney Extension to define well-formed sets in the case of the standard trimmed-NURBS representation [1, 4].

The data Δ in the trimmed-NURBS representation of a solid S is in two parts, the *geometric data* and the *topological data*. The geometric data in Δ comprises a finite set of compact oriented 2-manifolds-with-boundary, and corresponding sets of explicit boundary curves and corner vertices. The underlying surface for each 2-manifold is represented by a spline function over a parametric domain D_0 , with range in R^3 . The spline function is then restricted to a subset D of D_0 , as delineated by certain curves in the parametric domain, yielding a *trimmed NURBS patch*. (Such trimming may arise, for example, as the result of the intersection of two surfaces for representing the boundary of the solid.) Thus, each 2-manifold-with-boundary is represented by a trimmed NURBS patch.

Figure 1 illustrates two trimmed patches which join (approximately) along the intersection of two surfaces F and F' , restricted respectively to D and D' . The part of the representation of the solid corresponding to the intended in-

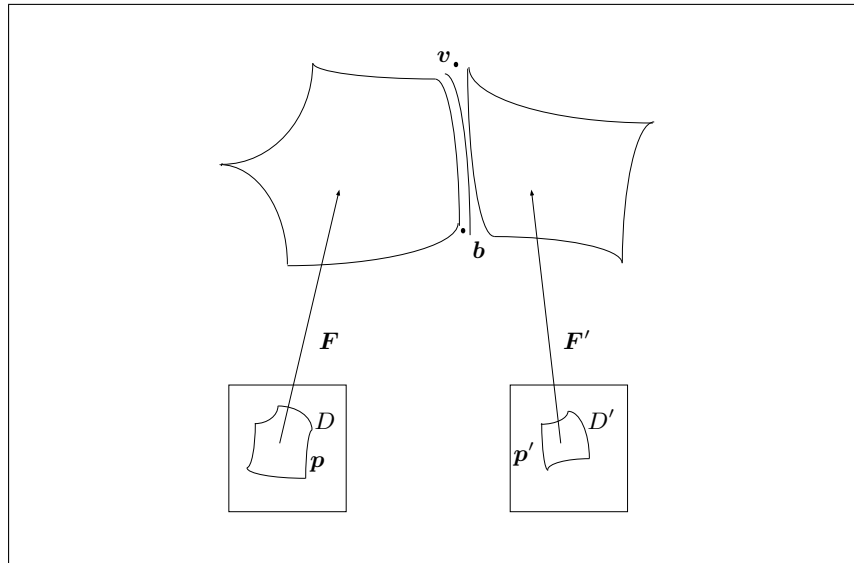


Fig. 1: Two adjoining trimmed NURBS patches

tersection comprises two (almost certainly inconsistent) pre-images, defined by *p-curves*, which are parametric curves p and p' with ranges in the respective parametric domains. In addition, there is often a third representation, usually

inconsistent with the other two, which is a parametric curve \mathbf{b} , with range in R^3 , and which follows closely the images of the p-curves \mathbf{p} and \mathbf{p}' . Finally, there may also be an explicit representation of each endpoint $\mathbf{v} \in R^3$ of the parametric curve $\mathbf{b} = \mathbf{b}(t)$, as illustrated in Figure 1. Here, however, we consider the more usual case where \mathbf{v} always coincides with an endpoint of \mathbf{b} .

The representation Δ also contains symbolic information, or *topological data*, describing how the faces, edges and vertices of the cellular decomposition of the boundary ∂S of S fit together. This data defines a topological 2-cycle. Ideally, the geometric and topological data are consistent: for example, corresponding to each 2-cell in the topological data is a trimmed surface patch in R^3 (two of these, $\mathbf{F}[D] = \{\mathbf{F}(u, v) : (u, v) \in D \subseteq D_0\}$ and $\mathbf{F}'[D'] = \{\mathbf{F}'(u, v) : (u, v) \in D' \subseteq D_0\}$, are shown in Figure 1).

Unfortunately, as illustrated in Figure 1, the curve $\mathbf{b}(t)$ does not usually coincide exactly with the corresponding edge of $\mathbf{F}[D]$, nor with the corresponding edge of $\mathbf{F}'[D']$, and the question therefore arises, which subset of R^3 should be considered to be represented by the given inconsistent data?

In [1], the set considered to actually be represented by Δ , *i.e.*, the *realization* S of Δ , is defined by its boundary ∂S . This boundary is made up of slightly perturbed trimmed-NURBS patches from Δ , where the perturbation is defined by the Whitney Extension Theorem. The slightly perturbed patches are not necessarily NURBS patches, but they are mutually consistent with the explicit boundary curves $\mathbf{b}(t)$, and they all fit together in a way that is exactly consistent with the topological data. This is done by taking $C = \partial D$ and

$$\epsilon(\mathbf{p}) = \mathbf{b}^k([\mathbf{p}^k]^{-1}(\mathbf{p})) - \mathbf{F}(\mathbf{p}), \quad \mathbf{p} \in \partial D$$

in Theorem 1, where $[\mathbf{p}^k]^{-1}$ is the inverse of the particular p-curve \mathbf{p}^k for which $\mathbf{p}^k(t) = \mathbf{p}$ for some $t \in [0, 1]$. Thus, ϵ is the difference between the edge of the trimmed patch $\mathbf{F}[D]$, and the boundary curves \mathbf{b}^k , viewed as a function of $\mathbf{p} \in \partial D$. (It is assumed that each \mathbf{b}^k and each \mathbf{p}^k is injective, and that distinct p-curves do not intersect, except at appropriate endpoints.) Then, the Whitney theorem can be used to extend the patch to all of D .

The meaning of the words ‘slightly perturbed’, in the previous paragraph, is thus quite satisfactory. The perturbation of the trimmed patch, denoted $\epsilon(\mathbf{p})$, is continuous, and the magnitude of the perturbation nowhere exceeds the magnitude of the largest discrepancy between the edge of the given patch and the given neighbouring explicit boundary curves $\mathbf{b}(t)$. (In other words, the perturbation is nowhere larger than the largest discrepancy already in the given data, along the edges of the given patch.) In addition, the perturbation satisfies a Lipschitz condition throughout the patch, with a Lipschitz constant for each component ϵ of ϵ equal to

$$L = \sup_{\mathbf{q}_1, \mathbf{q}_2 \in \partial D, \mathbf{q}_1 \neq \mathbf{q}_2} \frac{|\epsilon(\mathbf{q}_1) - \epsilon(\mathbf{q}_2)|}{\|\mathbf{q}_1 - \mathbf{q}_2\|} .$$

(In other words, the perturbation satisfies a Lipschitz condition over the *entire* patch, with a constant equal to the constant in the Lipschitz condition corresponding to the discrepancy present, along the edges of the given patch, in

the given data.) This Lipschitz condition is important, because it allows us to bound the change in the normal vector of the perturbed patch, relative to the normal vector \mathbf{n} of the given trimmed patch. Such a bound will be necessary when we want to preclude the possibility that adjacent patches have extraneous intersections. This question will be discussed in Section 4, below.

3.2 Combined-Subdivision-with-Trimming case

A representation permitting trimming of subdivision surfaces was proposed in [3], as an alternative to the classical NURBS representation. The representation is based on the combined subdivision schemes proposed by Levin [13], which permit the construction of subdivision surfaces having arbitrary boundary curves (any piecewise-smooth parametric curve possessing an evaluation procedure). Such schemes modify the subdivision stencils (for example, the Loop or Catmull-Clark stencils) near the boundary, using data from a boundary curve $\mathbf{c} = \mathbf{c}(u)$ supplied as part of the input [3].

We will refer to the representation of [3] as Combined Subdivision with Trimming (CST). Comparing it with the trimmed-NURBS representation of Subsection 3.1, the supplied boundary curve $\mathbf{c} : [0, 1] \rightarrow R^3$ corresponds to the combined closed curve made up from the m given boundary segments $\mathbf{b}^1(t), \dots, \mathbf{b}^m(t)$, joined end to end to form a simple closed curve embedded in R^3 . Since \mathbf{c} is required only to be piecewise smooth, it is possible to have corner points like those occurring at a join such as $\mathbf{b}^{l-1}(1) = \mathbf{b}^l(0)$ in the trimmed-NURBS representation. Beyond this, however, the representations are quite different. The surface patch in the CST representation is the limit of a modified combined-subdivision scheme which approximates an initially given input surface, and which, in the absence of roundoff error, interpolates the boundary curve $\mathbf{c} = \mathbf{c}(u)$ exactly. Note that the transfinite interpolation implied by this limiting process forms part of the actual representation. This is in contrast to the Whitney-Extension transfinite interpolation used in Subsection 3.1 for the purpose of defining a theoretical set determined by inconsistent data in a representation. We will, however, *also* apply Whitney-Extension transfinite interpolation in the case of CST representations.

Following [3], a general subdivision-surface control point i , at level j in the subdivision hierarchy, is denoted \mathbf{p}_i^j , and the breakpoint values in the parametric domain of \mathbf{c} are denoted u_i^j . A surface called the *original surface* [3], denoted here by Σ , is defined by the (given) control points \mathbf{p}_i^0 [3], along with a given subdivision scheme such as the Loop scheme [3, 13]. The subsequent trimming algorithm, which produces the actual trimmed-surface Σ_T , involves, first, a local remeshing of the control polyhedron to accommodate the trim curves, together with a sampling of Σ to choose control points for the trimmed surface. This information is used in an approximation stage that modifies the surface shape near the trim curve, in order to ensure proximity to the surface Σ .

The local remeshing will normally cause the assigned initial parameter values u_i^0 to be modified. Again following [3], we will continue to denote these modified

values by u_i^0 . The surface Σ is then sampled to find points corresponding to the control points of the original mesh.

The approximation stage of the trimming algorithm fits the trimmed surface Σ_T to Σ . This operation is only required near the trim curve, where the control mesh was generated by the remeshing algorithm. Away from this region, the trimmed surface is identical to Σ ; within the region, the approximation of Σ is generated as a hierarchy of *detail coefficients* \mathbf{d}_i^j , which are (additive) modifications of the control points \mathbf{p}_i^j . The depth of the subdivision hierarchy is limited by the introduction of a finite convergence threshold [3, Sec. 3.3].

The final trimmed-surface representation is therefore defined by a hierarchy of u -domain control points u_i^j , which depends on (the possibly modified versions of) the original breakpoints u_i^0 , and a hierarchy of control points \mathbf{p}_i^j with associated detail vectors \mathbf{d}_i^j , which depends on the original \mathbf{p}_i^0 . As observed in [3], the above trimming algorithm guarantees exact transfinite interpolation of the desired trimmed curve. This statement assumes that the subdivision process is continued until convergence, and without roundoff error. In practice, however, the actual representation will be the result of only a finite number of steps of the combined-subdivision process, and it will be the result of calculations using finite-precision floating-point arithmetic.

It follows from this last remark that if we place ourselves in the context of Subsection 3.1, and seek to prove rigorous mathematical theorems stating that the representation defines a well-formed subset of R^3 which has a boundary that is in some sense close to a given collection of input surfaces $\{\Sigma^k\}_{k=1,\dots}$, then we must agree on the definition of this subset. One possibility is to define the set in terms of the exactly defined limiting surfaces specifying its boundary patches. To do this, however, it would be necessary to preclude extraneous intersections, as defined at the end of Section 1: intersections between two boundary patches other than prescribed intersections along a common edge or at a common vertex. This, in turn, would require bounds on the normals of the limit surface Σ_T , which is defined in terms of the (rather complicated) process described in [3]. It is not clear whether this approach is feasible.

In this paper we propose an alternative approach, analogous to Subsection 3.1. The situation is illustrated in Figure 2, where two CST patches are intended to meet along a common edge. We assume, of course, that $\mathbf{c}^k(u)$ and $\mathbf{c}^{k'}(u)$ coincide along this common edge, but in practice, as already mentioned, the trimming algorithm will terminate after a finite number of steps and, furthermore, $\mathbf{c}^k(u)$ will be evaluated using finite-precision floating-point arithmetic; consequently, the edge of Σ_*^k , the computed approximation of Σ_T^k , will not coincide exactly with $\mathbf{c}^k(u)$. On the other hand (and this is implicit in [3]), it will be possible to obtain very satisfactory bounds for the difference, which we denote by $\epsilon(\mathbf{p})$, where \mathbf{p} is a point on the boundary of the final control polyhedron produced by the combined subdivision process. Assuming such bounds have been found, the Whitney Extension Theorem can be applied exactly as in Subsection 3.1 (the set C in Theorem 1 is now taken to be the boundary of the final control polyhedron) to provide a definition of the boundary of the actual




Fig. 2: Two adjoining CST patches

subset of R^3 specified by the inconsistent data that specified each of Σ_*^k and $\Sigma_*^{k'}$ separately.

The only remaining issue is the possibility of self-intersection of the combined patches: the patches provided by Whitney Extension guarantee that adjoining patches will meet properly at a common edge, but we must exclude the possibility that there are other intersections (see Figure 3, where the CST patches are shown in cross-section). In the case of CST patches based on Loop or Catmull-Clark subdivision, it has been shown [14] that Σ_*^k (respectively $\Sigma_*^{k'}$) and its derivatives can be evaluated directly, which means that the normal vector \mathbf{n}^k (respectively $\mathbf{n}^{k'}$) can be estimated, and, say, the method of Volino-Thalmann [15–17] applied. This is the subject of Section 4, where an analysis will be presented in a setting that includes both the trimmed-NURBS and CST cases.

4 Bounds on normal vectors

The fundamental step required, in order to apply the Volino-Thalmann method for detection of self-intersection of the boundary of the set defined by Whitney Extension, is to bound the normal vector of the perturbed surface in terms of the normal of the original surface. The well-formed set that we are attempting to define is specified in terms of the perturbed surface provided by Whitney Extension, but it is only the normal vector of the original surface that is available to us.

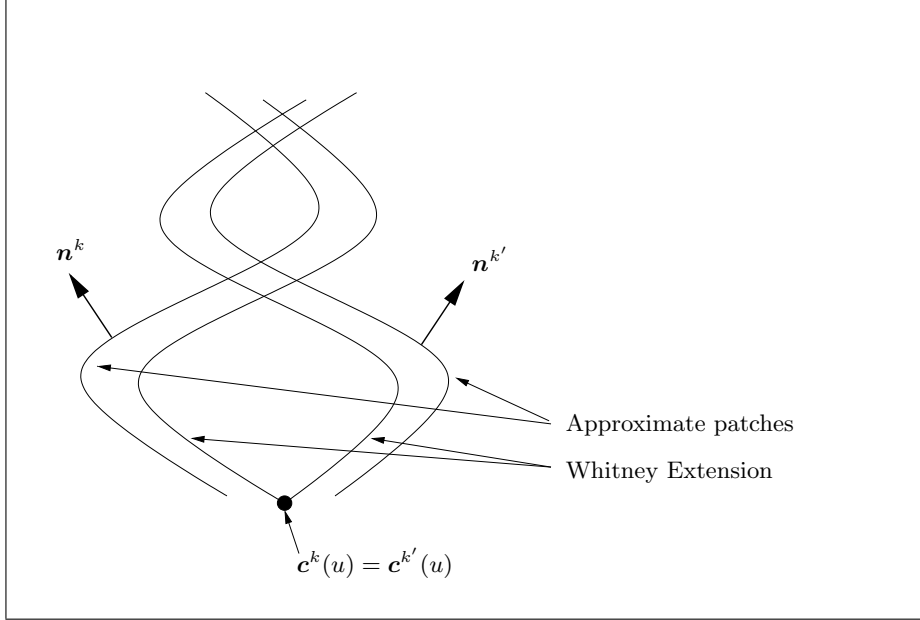


Fig. 3: Extraneous intersection of adjoining CST patches

The following analysis is an extension of that found in [1, Sec. 5].

Denote the original patch by \mathbf{F} , and the perturbed patch defined by Whitney Extension by $\mathbf{G} = \mathbf{F} + \epsilon$. (In the trimmed-NURBS case, the notation \mathbf{F} here refers to $\mathbf{F}[D]$, the restriction of the NURBS mapping $\mathbf{F}(u, v)$ to the trimming domain D , as in Subsection 3.1. In the CST case, Subsection 3.2, \mathbf{F} refers to a CST patch such as Σ_T^k .) Also, denote by \mathbf{n}_F and \mathbf{n}_G the normal vectors at an arbitrary parameter point (U, V) . (In the trimmed-NURBS case, see [18]; in the Loop or Catmull-Clark case, see [14].) In order to use a variational argument, we will write $U = u_o + \delta u$, $V = v_o + \delta v$, and consider the case when $U \rightarrow u_o$, $V \rightarrow v_o$. In the limit, the projection of $\mathbf{F}(u_o + \delta u, v_o + \delta v)$ on the tangent plane of \mathbf{F} at (u_o, v_o) is equal to $\mathbf{F}(u_o, v_o)$, and the normal $\mathbf{n}_F(u_o + \delta u, v_o + \delta v)$ is equal to $\mathbf{n}(u_o, v_o)$.

We will now give a bound for the angle ϕ , between the two normals \mathbf{n}_F and \mathbf{n}_G , expressed in terms of the coefficients

$$\begin{aligned} E &= \nabla_u \mathbf{F} \cdot \nabla_u \mathbf{F} \\ F &= \nabla_u \mathbf{F} \cdot \nabla_v \mathbf{F} \\ G &= \nabla_v \mathbf{F} \cdot \nabla_v \mathbf{F} \end{aligned}$$

of the first fundamental form:

$$I(du, dv) = Edu^2 + 2Fdu dv + Gdv^2$$

(see for example [19]).

Let $\mathbf{f} = \nabla_u \mathbf{F}(u_o, v_o)$ and $\mathbf{s} = \mathbf{n}(u_o, v_o) \times \nabla_u \mathbf{F}(u_o, v_o)$. In the orthogonal coordinate system $\left(\frac{\mathbf{f}}{\|\mathbf{f}\|}, \frac{\mathbf{s}}{\|\mathbf{s}\|}\right)$ of the tangent plane of \mathbf{F} at (u_o, v_o) , the coordinates of $\mathbf{F}(u_o + \delta u, v_o + \delta v)$ are:

$$\begin{aligned} u &= \mathbf{F}(u_o + \delta u, v_o + \delta v) \cdot \frac{\mathbf{f}}{\|\mathbf{f}\|} \\ v &= \mathbf{F}(u_o + \delta u, v_o + \delta v) \cdot \frac{\mathbf{s}}{\|\mathbf{s}\|}. \end{aligned}$$

We take u and v as new parameters, and, using the Implicit Function Theorem, define

$$\mathbf{m}(u, v) = \begin{pmatrix} u_o + \delta u \\ v_o + \delta v \end{pmatrix}.$$

Then, using Taylor's theorem, it is easy to show that

$$\mathbf{m}(u, v) = \begin{pmatrix} \frac{1}{\|\nabla_u \mathbf{F}\|} & \frac{-\cos \theta}{\sin \theta \cdot \|\nabla_u \mathbf{F}\|} \\ 0 & \frac{1}{\sin \theta \cdot \|\nabla_v \mathbf{F}\|} \end{pmatrix} \begin{pmatrix} u - \mathbf{F}(u_o, v_o) \cdot \frac{\mathbf{f}}{\|\mathbf{f}\|} \\ v - \mathbf{F}(u_o, v_o) \cdot \frac{\mathbf{s}}{\|\mathbf{s}\|} \end{pmatrix} + \begin{pmatrix} u_o \\ v_o \end{pmatrix} \quad (6)$$

where

$$\theta = \angle(\nabla_u \mathbf{F}(u_o, v_o), \nabla_v \mathbf{F}(u_o, v_o)) \quad (7)$$

and

$$\mathbf{G}(u, v) = (u, v, 0)^T + (\boldsymbol{\epsilon}(\mathbf{m}(u, v)) \cdot \frac{\mathbf{f}}{\|\mathbf{f}\|}, \boldsymbol{\epsilon}(\mathbf{m}(u, v)) \cdot \frac{\mathbf{s}}{\|\mathbf{s}\|}, \boldsymbol{\epsilon}(\mathbf{m}(u, v)) \cdot \frac{\mathbf{n}}{\|\mathbf{n}\|})^T. \quad (8)$$

The second term on the right is the perturbation vector expressed in the local coordinate system $\left(\frac{\mathbf{f}}{\|\mathbf{f}\|}, \frac{\mathbf{s}}{\|\mathbf{s}\|}, \frac{\mathbf{n}}{\|\mathbf{n}\|}\right)$.

To simplify the notation (and to emphasize the correspondence with the derivation in [1]), let $(e_1, e_2, e_3)^T$ denote the components of $\boldsymbol{\epsilon}$ in the parametric domain (u, v) . Then, from (8),

$$\mathbf{G}(u, v) = \begin{pmatrix} u + e_1 \\ v + e_2 \\ e_3 \end{pmatrix},$$

and in the limit as $\delta u, \delta v \rightarrow 0$,

$$\mathbf{n}_{\mathbf{F}}(u, v) = (0, 0, 1)^T.$$

It follows that if ϕ is the angle between $\mathbf{n}_{\mathbf{F}}$ and $\mathbf{n}_{\mathbf{G}}$, then in the limit

$$\cos \phi = \frac{(\mathbf{G}_u \times \mathbf{G}_v) \cdot (0, 0, 1)^T}{\|\mathbf{G}_u \times \mathbf{G}_v\|}$$

where

$$\mathbf{G}_u \times \mathbf{G}_v = \begin{pmatrix} 1 + e_{1u} \\ e_{2u} \\ e_{3u} \end{pmatrix} \times \begin{pmatrix} e_{1v} \\ 1 + e_{2v} \\ e_{3v} \end{pmatrix} = \begin{pmatrix} e_{2u}e_{3v} - e_{3u}e_{2v} - e_{3u} \\ e_{3u}e_{1v} - e_{3v}e_{1u} - e_{3v} \\ 1 + e_{1u} + e_{2v} + e_{1u}e_{2v} - e_{2u}e_{1v} \end{pmatrix}.$$

It can then be shown, using the same analysis as that in [1], that

$$\cos \phi \cong 1 - \frac{(e_{3u})^2}{2} - \frac{(e_{3v})^2}{2}. \quad (9)$$

In order to bound the derivatives e_{3u} and e_{3v} , we use (6) and the fact that from the Whitney Extension Theorem, if ϵ satisfies a Lipschitz condition with constant L on the boundary of the patch, then it satisfies the same Lipschitz condition everywhere, and consequently

$$\left| \frac{\partial e_3}{\partial m_i} \right| \leq L, \quad i = 1, 2.$$

Using this inequality, and (7) and (9), it follows (after some algebra) that

$$\cos \phi \geq 1 - \frac{L^2}{2} \left[\frac{E + 2F + G}{EG - F^2} \right]$$

where E , F and G are the coefficients of the first fundamental form, given above.

Note that if $\theta = \frac{\pi}{2}$, then $F = 0$, and this lower bound becomes $1 - \frac{L^2}{2} \left[\frac{E+G}{EG} \right]$. If in addition we have $\|\nabla_u \mathbf{F}\|^2 = \|\nabla_v \mathbf{F}\|^2 = 1$, so that $E = G = 1$, then the lower bound becomes $1 - L^2$. This was the special case treated in [1, Sec. 5].

It would be worthwhile, also, to study the variation of the normals (between the surfaces Σ and Σ_T) induced by the CST scheme, and to modify the scheme, if necessary, to reduce this variation. The additional variation introduced by Whitney Extension will normally be small. Indeed, the value of L used for Whitney Extension will be small provided that the curvature of \mathbf{c} is bounded below, and provided that the spacing u_i^j is fine enough, relative to this lower bound. (We are assuming here, however, that the spacing of the u_i^j is still coarse enough so that roundoff error does not play a major role in the variation of ϵ .)

5 Conclusion

In this paper we have shown how Whitney Extension can be used in the well-definition of objects defined by approximate surface patches, in conjunction with the bounds derived here on the variation of normal vectors. These bounds depend on the Lipschitz constant defined by the error in the approximate surface patch along its boundary, and by the (first-normal-form) parameters of the surface.

6 Acknowledgment

The authors are grateful to Ian Stewart, and to an anonymous referee, for several useful comments. Responsibility for errors or omissions rests solely with the authors.

References

1. Andersson, L.-E., Stewart, N. F. and Zidani, M. Error analysis for operations in solid modeling in the presence of uncertainty. Available at www.iro.umontreal.ca/~stewart.
2. Whitney, H. Analytic extensions of differentiable functions defined in closed sets. *Trans. Amer. Math. Soc.* (36), 63-89, 1934.
3. Litke, N., Levin, A. and Schröder, P. Trimming for subdivision surfaces. *Computer Aided Geometric Design* (18), 463-481, 2001.
4. Hoffmann, C. M. and Stewart, N. F. Accuracy and semantics in shape-interrogation applications. *Graphical Models* 67, No. 5, 373-389, Sept. 2005.
5. Gould, S. H. Variational Methods for Eigenvalue Problems. University of Toronto Press, 1957.
6. Sewell, M. J. Maximum and Minimum Principles. Cambridge, 1987.
7. Coons, S. A. Surfaces for Computer Aided Design of Space Forms. MIT Project Mac, TR-41, MIT, Cambridge, MA, June 1967.
8. Nielson, G. M. A transfinite, visually continuous, triangular interpolant. In Geometric Modeling: Algorithms and New Trends. G. E. Farin, ed. SIAM, 1987.
9. Gross, L. and Farin, G. A transfinite form of Sibson's interpolant. *Discrete Applied Mathematics* (93), 33-50, 1999.
10. Biswas, A., Shapiro, V. and Tsukanov, I. Heterogeneous material modeling with distance fields. *Computer Aided Geometric Design* (21), 215-242, 2004.
11. Aronsson, G. Extension of functions satisfying Lipschitz conditions. *Arkiv för Matematik* (6), No. 28, Stockholm, 1967.
12. McShane, E. J. Extension of range of functions. *Bull. Amer. Math. Soc.* (40), 837-842, 1934.
13. Levin, A. Combined subdivision schemes for the design of surfaces satisfying boundary conditions. *Computer Aided Geometric Design* (16), 345-354, 1999.
14. Stam, J. Exact evaluation of Catmull-Clark subdivision surfaces at arbitrary parameter values. *Proc. ACM SIGGRAPH*, 395-404, 1998.
15. Volino, P. and Thalmann, N. M. Efficient self-collision detection on smoothly discretized surface animations using geometrical shape regularity. *Eurographics '94* (13), No. 3, pp. C-155 to C-164. M. Daehlen and K. Kjelldahl, Guest Editors, Blackwell Publishers, 1994.
16. Grinspun, E. and Schröder, P. Normal bounds for subdivision-surface interference detection. Proceedings of the IEEE Conference on Visualization, 333-340, 2001.
17. Andersson, L.-E., Stewart, N. F. and Zidani, M. Proof of a non-selfintersection conjecture. To appear, *Computer Aided Geometric Design*. Available at www.iro.umontreal.ca/~stewart.
18. Piegl, L. and Tiller, W. The NURBS Book. Springer, 1997.
19. Pressley, H. Elementary Differential Geometry, Springer, 2001.