C^2 surfaces built from zero sets of the 7-direction box spline

Jörg Peters

Abstract. Curvature continuous surfaces can be generated by tracing the zero set of a trivariate box spline of degree four. This box spline is defined by seven directions that form a regular partition of space into tetrahedra so that an efficient approximate evaluation is obtained via subdivision and piecewise linear approximation. There is a natural representation of hierarchical data by perturbation at various levels of subdivision. The approach is contrasted with curvature continuous parametric splines.

§1. Introduction

Parametric curvature continuous surfaces following the outlines of a regular, checkerboard structured mesh are in common use in the form of bicubic tensor product splines. Accomodating a non checkerboard layout is far more difficult, because the problem is truly bivariate in nature. Only a few C^2 surface constructions for irregular layouts are currently known ([Seidel], [Hagen,Pottmann], [Hahn], [Peters]). None of these representations is completely satisfactory, placing restrictions on the input data, are slow to evaluate, lack guaranteed shape properties or are based on lenghty and hence not elegant formulas.

In contrast, it is straightforward to *define* a curvature continuous surface as the zero set of a trivariate, piecewise polynomial, C^2 function. The zero set of such a function is of the same smoothness class since the expansions of the derivatives agree in all directions. By choosing the coefficients of this function suitably, we may hope to capture the geometry and allow for an approximate parametrization of the surface. There are currently two main approaches to defining the individual pieces of a function whose zero set represents a surface. The first is to locally generate a shell-like structure consisting of polygonal cells that follow the outlines of and enclose the intended surface. This approach is illustrated by the recent work in [Sederberg], [Dahmen], [Guo], [Dahmen, Thamm-Schaar], [Bajaj] and [Middleditch]. Creating the appropriate shell structure can be challenging and leads to special cases, say in the case of coplanar pieces. The alternative approach defines a function on a regular,

IMA Conference on Mathematics of Surfaces VI
R.R. Martin, G. Mullineux, M.A. Sabin (eds.), pp. 1–5.
ISBN xxx.

global lattice. This has the advantage that non checkerboard features of the surface do not require special treatment, but is off-hand less efficient, because n^3 lattice values have to be stored, rather than just $O(n^2)$ plus the shell structure in the case of a surface following implicit approach. However, the efficiency of the global approach can be improved by maintaining only the lattice points close to the intended surface; that is, the surface is enclosed in a rectilinear bounding volume.

The real zero set of a trivariate function consist in general of several connected components. The components may touch one another yielding a non-unique surface continuation, and they may contain handles and holes that do not reflect the design intent [Guo]. A guarantee of existence and uniqueness, i.e the identification of a singly connected sheet in the region of interest, is therefore a main concern when choosing an implicit surface representation. The algebraic degree of the surface pieces is of interest in this context, because a higher degree allows for more complex zero sheets. This makes it in general more difficult to identify the zero sheet that represents the surface and to prove properties of the resulting surface.

Considering regular lattices, there are a number of options for constructing piecewise polynomial C^2 functions. Maybe the simplest choice is to generalize the cubic B-spline construction by choosing a piecewise tricubic, tensor-product function. This approach is unsatisfactory for two reasons. First, for any tensor-product function there is ambiguity when choosing a simple piecewise linear surface approximation based on the values at the vertices of the lattice. The natural process for obtaining an approximation in the context of splines is to choose the values to be the coefficients of the spline and the cube-like lattice cells to be the regions were the piecewise polynomial is infinitely differentiable; approximate evaluation to high precision can then be achieved by subdivision. However, specifying values at the vertices of a cube and invoking the Rolle's theorem for continuous functions to obtain the connectivity of the zero sheet is not a well-defined operation. Already looking at a face of the cube with the sign pattern $\pm\mp$ allows for two different piecewise linear zero sheet approximations attached to the midpoints of edges with changing sign pattern: one with two sheets running from the upper left to the lower right and the other with two sheets running from the upper right to the lower left. For uniqueness of the zero sheet, it is desirable to have tetrahedral regions associated with the regular mesh, because a tetrahedron defines a unique connectivity of the midpoint based zero sheet for any given plus-minus pattern at the vertices: the least degree zero sheet separating positive and negative values at the midpoint of the corresponding edge is either empty, a triangle or a bilinear quadrilateral patch.

A second disadvantage of piecewise tricubics is their high total polynomial degree, 9. In contrast, the 7-direction box spline to defined below is of total degree 4 and is defined over tetrahedra. Simple averaging of the coefficients of the box spline attached to a regular lattice that partitions space allows for a stable approximate evaluation to high accuracy.

The paper is structured as follows. In the second section, the class of box splines to which the 7-direction box spline belongs is characterized. Then the 7-direction box spline is defined and the method for finding the zero set is sketched. The analysis concludes with a discussion of the properties of the box spline zero set and a comparison with parametric curvature continuous spline surfaces.

§2. Box splines with unit directions

Box splines are defined by a set of n directions which determine both the support of the piecewise polynomials and their continuity properties. For the purpose of this paper it suffices to look at box splines whose direction set contains all unit vectors of the domain of the box spline. Examples of this class of box splines are all univariate B-splines over uniform knot sequences, the bivariate Zwart-Powell element and of course the 7-direction box spline which is the subject of this paper. For a detailed discussion of box splines in a general setting I find it useful to consult [de Boor, Höllig, Riemenschneider].

We define the box spline with respect to the m unit vectors of $\mathbb{R}^m$ as the characteristic function of the unit cube spanned by the vectors. This box spline has degree zero and is discontinuous. To obtain a box spline with a larger direction set, the lower order box spline is successively convolved in each of the remaining directions, thus increasing the degree and the support but not necessarily the continuity. The degree increases by one with each direction, and the support grows by forming the Minkowski sum of the previous support with the unit cube shifted in the direction. To determine the continuity of the resulting box spline we need to determine the number a, which counts the minimal number of directions that need to be removed from the direction set to obtain a reduced set that does not span $\mathbb{R}^m$. Then the continuity class is C^{a-2}.

Example 1. The univariate uniform cubic spline has the direction set $\{1, 1, 1, 1\}$ (cf. Figure 2.1, right). Hence $m = 1$, $n = 4$, and $a = 4$ since all elements of the set have to be removed to make it nonspanning. The degree is therefore $n - 1 = 3$ and the continuity is of order $a - 2 = 2$.

Fig. 2.1. Uniform univariate splines.

Example 2. The bivariate Zwart Powell element is the box spline with the direction set $\{(1,0),(0,1), (1,1),(-1,1)$. We have $m = 2$, $n = 4$, $a = 3$ and hence the degree of the element is 2 and the order of continuity is 1. Figure 2.2 shows the support of box splines for a given set of directions.

Example 3. The continuity of the box spline with the direction set $\{(1,0),(1,.5),(1,1),(.5,1),(0,1)\}$ is a maximal $4-2$. However, the resulting tesselation of the domain consists of both triangles and quadrilaterals which makes it difficult to represent the polynomial pieces in standard form.

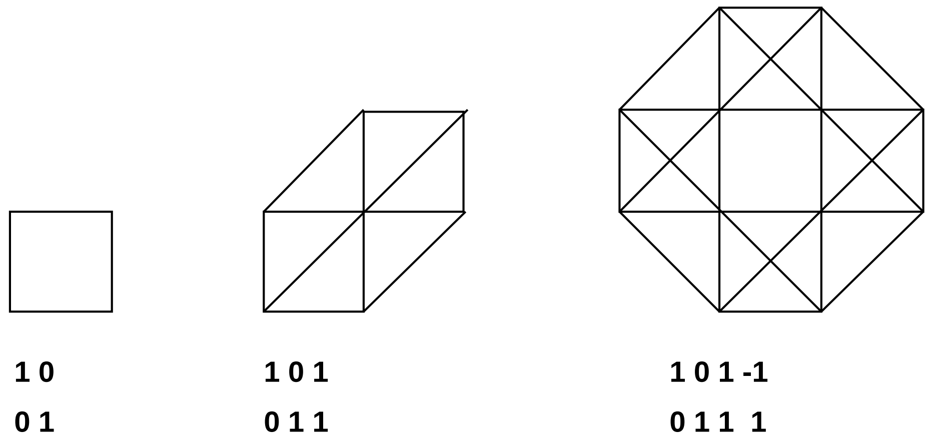

Fig. 2.2. The support regions of bivariate box. splines with the given direction set

§3. The 7-direction box spline

The Zwart-Powell element is special among the low degree box splines defined over the plane, in that it has maximal smoothness $n-1$ and is piecewise polynomial over a regular triangulation. The 7-direction box spline is a similar serendipity element among the trivariate box splines. The seven directions

$$\begin{array}{ccccccc} 1 & 0 & 0 & 1 & -1 & 1 & -1 \\ 0, & 1, & 0, & 1, & 1, & -1, & -1 \\ 0 & 0 & 1 & 1 & 1 & 1 & 1 \end{array}$$

cut $\mathbb{R}^3$ into a symmetric regular arrangement of tetrahedra.

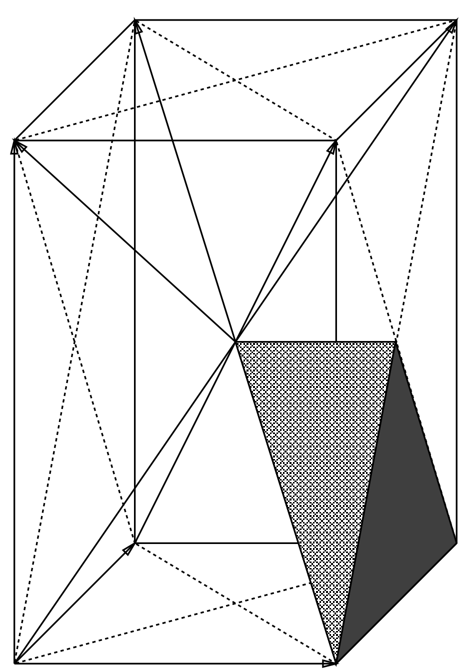

Fig. 3.1. The 7 directions of the box spline and its domain tetrahedra.

A simple count yields $m = 3$, $n = 7$ and $a = 4$. Thus the polynomial piece defined over each tetrahedron is of degree four and the spline formed as a linear combination of box splines is C^2. The support has the shape of an octahedron, consisting of 5 cubes along the diagonals.

To model with the zero set of the 7 direction spline, we partition space into cubes by a regular lattice. Each lattice point has a real number associated with it. That is, we create a *field* and this field defines defines the spline since we interpret its values as the coefficient values of the spline. For example, the lattice point value can be the signed distance to some piecewise linear outline of the smooth surface to be constructed (cf. [Wyvill]).

A unique *approximate evaluation* of the surface is achieved as follows. Consider the tesselation of each cube into tetrahedra and associate the average of the values at the vertices with the center of each cube or cube face. This allows the construction of a continuous piecewise linear approximation to the zero sheet. For each edge whose endpoints have an opposite sign, mark the midpoint. Each tetrahedron has either zero, three or four marked edge-midpoints. Correspondingly, we add no, one or two (coplanar) triangles connecting the midpoints to a list of triangles. The union of the triangles in the list then form the surface approximation.

The surface approximation is refined by averaging according to the subdivision rules of the 7-direction box spline. That is, each value is replicated over a cube of half the edge length and then the values on this refined lattice are averaged consecutively in each of the four diagonal directions of the box spline. The quadratic convergence of the coefficients of the box spline to the surface [de Boor, Höllig, Riemenschneider, (30)Theorem] assures us that the sequence of linear approximations converges quadratically to the surface. That is, we need not evaluate the box spline explicitly ever, but rather evaluate approximately up to a tolerance defined by the output requirements.

The evaluation by averaging implies that features are smoothed out at each step. This has the desirable effect that no additional features are introduced and hence the shape of the surface can be inferred from the first approximation. In particular, no additional zero sheets can be generated so that the surface is *single sheeted* if the first approximation is single sheeted. The undesirable effect of the averaging procedure is a washing out of features. This washing out can be regulated and even completely prevented by perturbing the values of the refined lattice at a step of the subdivision. Theoretically, if we perturb at a finite number of subdivision steps, we obtain in the limit a C^2 surface based on the finer lattice. This allows defining sharper features. If we perturb at each subdivision step, we may even reduce the continuity to model edges and corners. In other words, the pertubation at each step allows a hierarchical data representation.

§4. Discussion and comparison

In this section, the implicit C^2 surface generation via zero sets is compared to the parametric representation exemplified by the approach in [Peters]. Both the parametric and the implicit approach have demonstrated the capability to model free-form surfaces (cf. Figure 4.1, respectively [Peters] or http://www.cs.purdue.edu/people/jorg) However, the manipulation of the blend ratios in [Peters] provides a more natural set of *shape parameters* than the manipulation of coefficients at various stages of subdivision in the present approach.

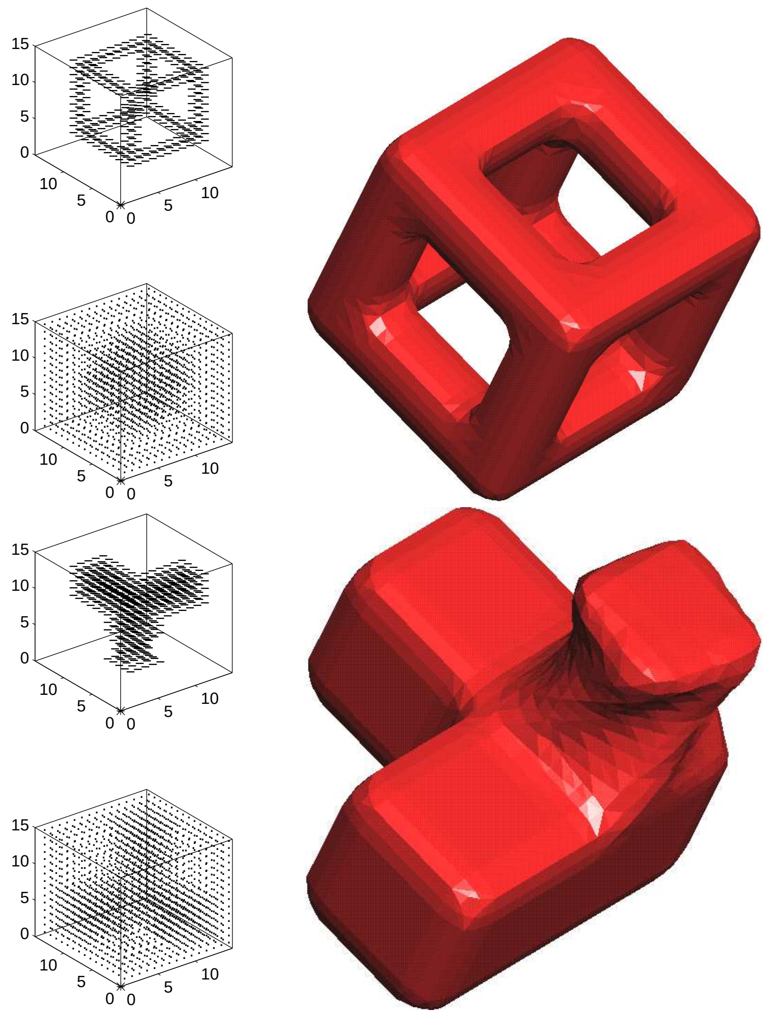

Fig. 4.1: Two free-form surfaces as zero sets of the 7-direction box spline: The figures on the left show the negative and positive field entries. The figures on the right show a (flat shaded) piecewise linear approximation to the surface.

An important property both of parametric and of implicit surface representations is the *piecewise linear approximability* to facilitate the location of points with respect to the surface, respectively to be able to render the surface. In fact a piecewise implicit surface representation may also require a piecewise linear approximation in order to test whether a point is inside or outside, since a checking of the pieces is not sufficient to establish this property globally (see Figure 4.2).

For parametric surfaces a refinable piecewise linear approximation in the form of the local convex hull is useful for a variety of tasks, for example inter-

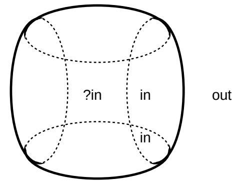

Fig. 4.2. Point classification for piecewise implicits may be difficult.

section tests. Here the local convex hull refers to the convex hull of the input data in a small neighborhood of a surface point. Remarkably few methods in the literature on free-form surface modeling techniques can guarantee their surface to lie in the local convex hull of the input data which may account for the oscillations observed in surfaces generated by methods where the property has not been established [Mann et al]. One reason why methods have not been tested for this important property is that proving this property for parametric surfaces is in general difficult: typically coefficients are derived from continuity constraints in terms of difference equations; showing the convex hull property requires them to be expressed as averages of input data. By contrast the field defining the 7-direction box spline naturally localizes points and defines inside and outside.

The hard part for implicitly defined surfaces is usually to obtain an approximate linear parametrization of the surface. For this *single-sheetedness* of the zero set in the region of interest is usually required. Single-sheetedness in the region of interest is usually trivial for parametric surfaces but requires a careful analysis and often nonlinear or inequality constraints on the coefficients of implicitly defined surfaces [Guo, Bajaj]. Consequently, the space of functions with single-sheeted zero-sets is generally not closed under basic operations such as averaging. As Figure 4.3 illustrates, the zero set of the sum of two functions each of which has a single zero sheet may have infinitely many zero sheets. Here we took advantage of the opposite curvatue of functions. However, even if the values are monotone say in the y-direction, Figure 4.4 shows that single-sheetedness may not be preserved. As a partial remedy, the evaluation of the 7-direction box spline by a piecewise linear approximation of the zero set defined by the coefficients forces locally a unique surface for each tetrahedron.

While the surface splines in [Peters] form a vector space for given connectivity of the defining mesh and blend ratios, the requirement of local single-sheetedness is nonlinear and is in general not preserved under addition of the coefficients of two box splines. In Figure 4.4 each column show the zero contour (computed by the Matlab contour command) of a function based on the four-direction box spline with increasing level of subdivision. The values above the graph in the two left columns are positive and negative below to avoid the anomaly depicted in Figure 4.4. However, the values are such that the zero-set of the sum of the functions is two-sheeted as displayed in the third column.

§5. Concluding Remarks

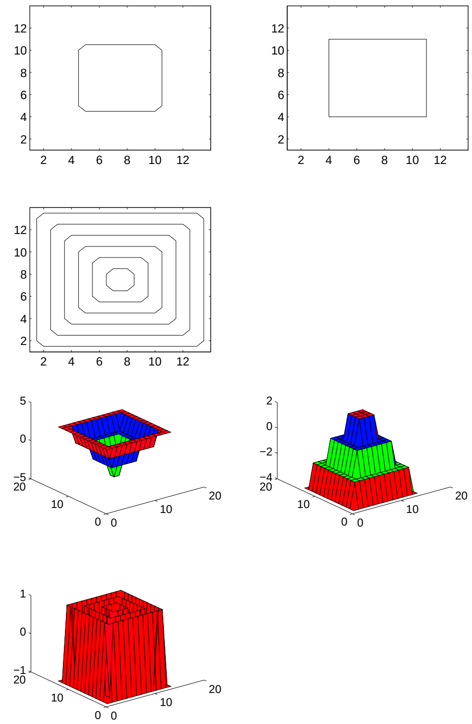

Figure 4.3: The zero set of of the sum of two functions each of which has a single zero sheet may have infinitely many zero sheets.

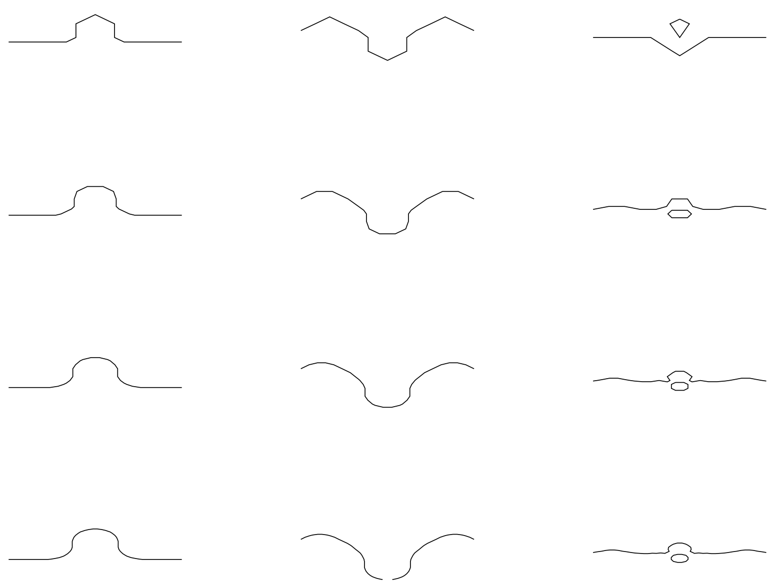

Figure 4.4: The sum of two locally single-sheeted box splines with monotonically increasing values over the 4-direction mesh is not necessarily single-sheeted.

The preceding sections have shown that the generation of smooth surfaces as the zero set of a piecewise polynomial function on a fixed-grid is conceptually simple and capable of modeling free-form objects. At present however, the simplicity of the approach is offset by a lack of efficiency. To be competitive with parametric methods the fixed-grid approach has to be supported by appropriate sparse data structures to reduce the complexity of representation. Such a representation has to use a consistent choice of field that guarantees that the average of two single sheeted implicits is again single-sheeted and facilitate generic graphics operations like rotation in a fixed grid. Further experimentation will show whether the simplicity of the approach translates into a viable modeling paradigm.

Acknowledgements: I thank Aaron Bilger for his help in generating Figure 4.1.

References

1. C. Bajaj, J. Chen, G. Xu, Interactive Modeling with A-patches, Computer Science CSD-TR-93-002, Purdue University, 1993.
2. J. Bloomenthal. Polygonalization of implicit surfaces, *Computer Aided Geometric Design*, 5(4):341–355, Nov 1988.
3. W. Dahmen, Smooth piecewise quadratic surfaces, In *Mathematical Methods in Computer Aided Geometric Design*, T. Lyche, L.L. Schumaker editors, Academic Press, Boston, 1989, pp 181–193.
4. W. Dahmen, T.-M. Thamm-Schaar, Cubicoids: modeling and visualization, *Computer Aided Geometric Design*, 10:93–108, 1993.
5. C. De Boor, K. Höllig, S. Riemenschneider, *Box splines* Springer Verlag, 1994
6. J.A. Gregory, *Smooth parametric surfaces and n-sided patches,* Computation of curves and Surfaces, W. Dahmen, M. Gasca and C.A. Micchelli, eds., Kluwer Academic Publishers, Dordrecht, 1990, pp. 457–498.
7. B. Guo, *Modeling Arbitrary smooth objects with algebraic surfaces* PhD thesis, Computer Science, Cornell University, 1991.
8. B. Guo, *Surface generation using implicit cubics* PhD thesis, Computer Science, Cornell University, 1991.
9. B. Guo, Non-splitting macro patches for implicit cubic spline surfaces, *Computer Graphics Forum*, 12(3):434–445, 1993.
10. H. Hagen, H. Pottmann, *Curvature continuous triangular interpolants*, in-Mathematical Methods in CAGD, T Lyche, L.L. Schumaker, eds. Boston Academic Press, 373–384.
11. J. M. Hahn, *Filling polygonal holes with rectangular patches,* Theory and Practice of geometric modeling, W. Straßer and H.-P. Seidel eds., Springer 1989.
12. Middleditch, Dimas, *Solid models with piecewise algebraic free -form faces,* CSG, April 1994.
13. J. Peters, Curvature continuous spline surfaces over irregular meshes, Technical Report, Purdue University, Department of Computer Sciences, CSD-TR-94-040, June 1994.
14. H. P. Seidel, *Symmetric recursive algorithms for surfaces: B-patches and the de Boor agorithm for polynomials over triangles*, Constructive Approximation 7(1991): 257–279.

15. T. Sederberg, Techniques for cubic algebraic surfaces, tutorial part ii, *IEEE Computer Graphics and applications*, 10(5):12–21, Sept 1990. *Implicit* , Siggraph
Warren, Aloha
16. B. Wyvill, C. Mc Pheeters, G. Wyvill, Animating soft objects, *Visual Computer*,

Acknowledgements. Jörg Peters is supported by NSF RIA 9396164-CCR and NSF NYI 9457806-CCR.

Jörg Peters
Department of Computer Science,
Purdue University,
W-Lafayette IN 47907-1398
USA
http://www.cs.purdue.edu/people/jorg