Eurographics Symposium on Geometry Processing 2010
Olga Sorkine and Bruno Lévy
(Guest Editors)

Volume 29 (2010), Number 5

Designing Quad-dominant Meshes with Planar Faces

Mirko Zadravec[1], Alexander Schiftner[2,3], and Johannes Wallner[1,2]

[1]Graz University of Technology, [2]Vienna University of Technology, [3] Evolute GmbH

Abstract

We study the combined problem of approximating a surface by a quad mesh (or quad-dominant mesh) which on the one hand has planar faces, and which on the other hand is aesthetically pleasing and has evenly spaced vertices. This work is motivated by applications in freeform architecture and leads to a discussion of fields of conjugate directions in surfaces, their singularities and indices, their optimization and their interactive modeling. The actual meshing is performed by means of a level set method which is capable of handling combinatorial singularities, and which can deal with planarity, smoothness, and spacing issues.

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [I.3.5]: Computational Geometry and Object Modeling—

1. Introduction.

Among the basic problems in freeform architecture which pertain to geometry processing is the decomposition of a freeform shape into a quad-dominant mesh whose faces are planar (or as planar as required by the intended manner of realizing that quad mesh in building construction). The geometric properties of such *PQ meshes* and an optimization procedure to generate them has been investigated by Liu et al. [LPW*06]. Pottmann et al. took this work further and discussed the *multilayer* constructions associated with such meshes [PLW*07]. It turns out that except for trivial cases, triangle meshes do not support offsets at constant distance or indeed any multilayer structure of meshes where corresponding faces are parallel. This again confirms the importance of quad meshes.

PQ meshes which are the basis of an architectural design must exhibit properties different from the mere geometric constraint of planarity of faces. A typical constraint is that faces are not larger than available panels of a given material, for example glass. At the same time face sizes should not vary much or be very small either, since that would lead to practical difficulties when realizing such a mesh as a structure. It is therefore important to incorporate the equal spacing of vertices in our design procedure.

Another issue is aesthetics, which we approach by, among others, the discrete bending energies of the polylines formed by the edges in a mesh. As already mentioned by [LPW*06],

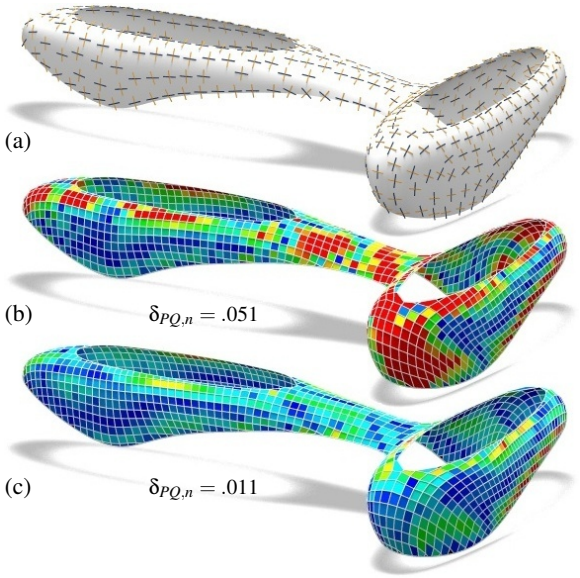

Figure 1: *Processing pipeline demonstrated for the outer hull of the Yas Island Marina Hotel, Abu Dhabi (Asymptote Architecture, originally a nonplanar quad mesh). (a) We start by optimizing a field of conjugate directions. (b) A level set method yields a quad mesh aligned with this field. (c) Optimization for planarity of faces. The color coding shows a normalized measure of planarity (normalized diagonal distance in quadrilateral faces; maximum value after optimization is $\delta_{PQ,n} = 0.011$; red color is used for values ≥ 0.01).*

the segmentation of a freeform surface Φ into quadrilateral planar faces is a discrete version of a *conjugate curve network* in Φ: optimization of that mesh towards planarity in general succeeds if it follows a network of conjugate curves, and cannot be expected to succeed otherwise. A similar result is true for segmentation of a surface into single-curved strips [PSB*08], which can be seen as a limit case of quad-meshing.

Contribution of the present paper. The papers mentioned above do not discuss the question how to *design* a network of conjugate curves, except to describe the degrees of freedom in principle, and to emphasize that only such networks are useful for meshing whose curves intersect transversely. The most important example of a conjugate curve network is the network of principal curvature lines: here the intersection angle of curves is 90 degrees and the transversality condition is fulfilled. In this paper we take up the design of *conjugate curve networks* and treat the following topics:

- The *conjugacy* relation between surface tangents, which is defined by the curvatures of the surface;
- an interpretation of the conjugacy relation in projective geometry terms which later allows us to encode a pair of conjugate directions by means of a single vector;
- fields of conjugate directions, and theoretical results on the elimination of singularities of such fields;
- a method based on level sets which converts direction fields into curve networks and subsequently into quad meshes or quad-dominant meshes.
- optimizing meshes generated in this way towards the PQ property.

This list of topics describes a processing pipeline for surfaces which we wish to represent as a PQ mesh. The last item is not a contribution of the present paper: we use a method similar to [LPW*06] for that.

Related Work. There are papers on quad meshing, and also on approximating a surface by a mesh with planar faces. For quad meshing in general we refer the reader to the literature cited in the introduction of [LPW*06]. Recent work which can create near-planar faces is [BZK09], where a cross field is aligned with dominant principal curvature lines, and subsequently a quad mesh is roughly aligned with the principal curvature lines by means of a continuous-discrete optimization problem. The examples given in that paper however make it clear that planarity of faces is not intended: the principal curvature directions are taken only as guidelines to achieve a good meshing. Exact planarity is achieved by [CAD04], but here smoothness is not an issue. These methods cannot be used for our purposes without modification.

In this paper singularities of direction fields play a prominent role; see [RVLL08, RVAL09] for directly related work. The present paper is different from previous ones on meshing because the constraint of planarity of faces is very rigid,

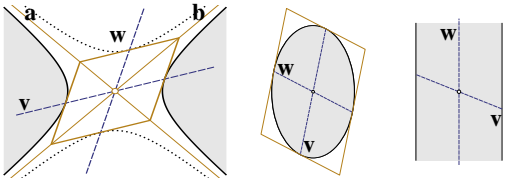

Figure 2: *The conjugacy relation* $\mathbf{v}^T M \mathbf{w} = 0$ *and the conics* $\mathbf{x}^T M \mathbf{x} = \pm 1$ *(Dupin indicatrices). If* $\det M \geq 0$*, only one of these conics is real. If* $\det M < 0$*, then* $\mathbf{v}, \mathbf{w}$ *are separated by self-conjugate vectors* $\mathbf{a}, \mathbf{b}$ *which indicate the asymptotes.*

and in our attempts at meshing a surface we have to exactly follow the features of that surface which are relevant to planarity. Many of these "second order" features would not be called features in other circumstances.

It must be mentioned that there exist methods for generating PQ meshes which entirely circumvent the problem of surface analysis: In order to approximate a desired shape Φ by a dense PQ mesh we can first approximate Φ by a coarse one, and subsequently apply several rounds of PQ optimization and subdivision in an alternating way [LPW*06]. But even if optimization is done such that deviation from Φ is penalized, the second order features of the original shape Φ will be different from those of the mesh actually obtained (this is clear from the later section on invariance of singularities, for an illustration see Figure 11).

2. Conjugate Directions in Surfaces.

This section introduces the notion of *conjugacy* of surface tangents which involves the surface's second derivatives and is defined as follows [dC76]: Any point $\mathbf{p}$ of a smooth surface Φ serves as the origin of a coordinate frame whose x_3 axis is orthogonal to Φ. Φ can locally be described as the graph of a function $x_3 = f(\mathbf{x})$, with $\mathbf{x} = \binom{x_1}{x_2}$, and which has the 2nd order Taylor expansion

$$x_3 = \frac{1}{2}\mathbf{x}^T M \mathbf{x} + \ldots, \text{ where } M = \begin{pmatrix} \partial_{11} f & \partial_{12} f \\ \partial_{12} f & \partial_{22} f \end{pmatrix}.$$

It is well known that the principal curvatures κ_1, κ_2 are the eigenvalues of M, and the principal directions correspond to the eigenvectors. The Gaussian curvature equals

$$\kappa_1 \kappa_2 = K = \det(M).$$

If tangent vectors $\mathbf{v}, \mathbf{w}$ are represented as elements of $\mathbb{R}^2$, then $\mathbf{v}^T M \mathbf{w}$ evaluates the second fundamental form for them. We define

$$\mathbf{v}, \mathbf{w} \text{ are conjugate } \iff \mathbf{v}^T M \mathbf{w} = 0.$$

Obviously, if $\mathbf{v}$ and $\mathbf{w}$ are conjugate then so are any nonzero multiples of these vectors and we can speak of *conjugate*

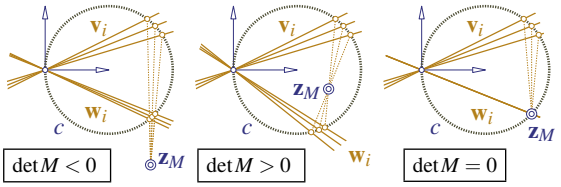

Figure 3: *Illustrating Prop. 1 for vectors* $\mathbf{v}_i$ *and corresponding conjugate vectors* $\mathbf{w}_i$*. If* $\mathbf{v}$ *rotates about the origin clockwise, then* $\mathbf{w}$ *rotates clockwise or counter-clockwise, depending on the sign of* $\det M$*.*

directions – a direction being the linear span of a tangent vector.

The set of vectors conjugate to $\mathbf{v}$ is a 1-dimensional subspace except in the special case that $\det M = 0$ and $M\mathbf{v} = 0$. We can visualize the conjugacy relation by means of the *Dupin indicatrices*, which are conics with the implicit equation $\{\mathbf{x} \mid \mathbf{x}^T M \mathbf{x} = \pm 1\}$ (see Figure 2).

The following elementary fact is well known in projective geometry, where it belongs to 'involutions on a conic' [Cox92, § 7.5]. Consider the auxiliary circle

$$c : (x_1 - 1)^2 + x_2^2 = 1$$

and for each line $[\mathbf{v}]$ spanned by a vector $\mathbf{v}$ consider its intersection point $[\mathbf{v}] \cap c$. Here we abuse notation and do not count the trivial intersection point $\mathbf{o} = (0,0)$, so $[\mathbf{v}] \cap c$ is always a single point. Then the following is true:

Prop. 1. *Vectors* $\mathbf{v}, \mathbf{w}$ *are conjugate, that is* $\mathbf{v}^T M \mathbf{w} = 0 \iff$ *the points* $[\mathbf{v}] \cap c$, $[\mathbf{w}] \cap c$ *and* $\mathbf{z}_M = \frac{2}{m_{11}+m_{22}} \begin{pmatrix} m_{22} \\ -m_{12} \end{pmatrix}$ *lie on a common straight line. The point* $\mathbf{z}_M$ *lies inside* c*, outside* c*, or on* c *if* $\det M > 0$*, or* $\det M < 0$*, or* $\det M = 0$*, respectively.*

This is illustrated by Figure 3. The point $\mathbf{z}_M$ may also escape to infinity, in which case the formula above reads $\mathbf{z}_M = \frac{1}{0} \begin{pmatrix} m_{22} \\ -m_{12} \end{pmatrix}$. Apparently $\mathbf{z}_M$ determines the matrix M up to a scalar factor.

Data structure for the conjugacy relation. The procedures described later in this paper require that we deal with curvatures and the conjugacy relation in all points of a surface Φ. This is implemented as follows: Φ is represented by a triangle mesh (V, E, F). We store the necessary curvature information in the faces, which are equipped with a local coordinate system.

We use the method of [FSDH07] to find a nonzero smooth "basis" vector field $\mathbf{b}$ (by prescribing a value in a face and minimizing the Dirichlet energy). It is represented as a discrete 1-form and is evaluated for each face $f \in F$. By normalization we get the first basis vector $\mathbf{e}_{1,f} = \mathbf{b}_f / \|\mathbf{b}_f\|$ of f's local coordinate system. The second basis vector then equals $\mathbf{e}_{2,f} = \mathbf{n}_f \times \mathbf{e}_{1,f}$, where $\mathbf{n}_f$ is the positive unit normal vector. Such a field $\mathbf{b}$ exists if Φ has disk topology, which is the case for the applications we have in mind. The matrix M_f describing the second fundamental form in f is found from its eigen-data (principal curvatures and principal directions), using the method of osculating jets [CP03]. We further compute the point $\mathbf{z}_{M,f}$ according to Prop. 1.

Remark: It is well known that principal directions are numerically unstable if curvatures are almost equal, but M_f is stable in any case; $\mathbf{z}_{M,f}$ is stable whenever $M_f \neq \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$.

3. Fields of transverse conjugate directions (TCD fields).

In our study of conjugate curve networks we encounter the problem of assigning, to each point $\mathbf{p}$ of a surface, a pair of conjugate directions. We could do this locally by assigning two vectors to each point which must obey the side condition of conjugacy, but globally this is usually not possible. Besides it is better for subsequent optimization tasks if we find a representation which avoids side conditions altogether. Our way of choosing a conjugate pair of directions is based on the matrix M which stores curvature information and an auxiliary matrix N which expresses our choice. Actually we work with the corresponding points $\mathbf{z}_M, \mathbf{z}_N$ instead of the matrices M, N. This approach is based on the following proposition:

Prop. 2. *If* 2×2 *matrices* M, N *are symmetric with* $\det N > 0$*, then there are nonzero vectors* $\mathbf{v}, \mathbf{w}$ *which fulfill* $\mathbf{v}^T M \mathbf{w} = \mathbf{v}^T N \mathbf{w} = 0$*. They are eigenvectors of the matrix* $N^{-1}M$*.*

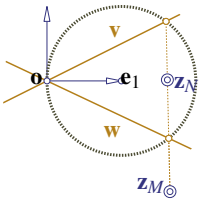

This is well known in linear algebra, but is also a consequence of Prop. 1: $\mathbf{v}, \mathbf{w}$ are found by intersecting the line $\mathbf{z}_M \mathbf{z}_N$ with the circle c. This intersection surely exists if $\mathbf{z}_N$ lies inside c. The line $\mathbf{z}_M \mathbf{z}_N$ is well defined also in case $\mathbf{z}_M$ is at infinity. If $\mathbf{z}_M$ is already known, then we can store a pair of conjugate vectors $\mathbf{v}$, $\mathbf{w}$ simply by storing $\mathbf{z}_N$.

The transversality condition. With Prop. 2 we can choose a pair of conjugate directions by choosing the point $\mathbf{z}_N$. In our implementation, we have such a point $\mathbf{z}_{N,f}$ for each face. We can thus view $\{\mathbf{z}_{N,f}\}_{f \in F}$ as a vector field. All future optimization problems regarding conjugate directions are converted into problems regarding the vector field $\mathbf{z}_N$. One example of how properties of conjugate directions are expressed in terms of the vector field $\mathbf{z}_N$ is the relation

$$\cos \sphericalangle(\mathbf{v}, \mathbf{w}) \leq \|\mathbf{z}_N - \mathbf{e}_1\|, \tag{1}$$

whose proof is elementary and follows from the fact that $2\sphericalangle(\mathbf{v}, \mathbf{w})$ occurs in the vertex $\mathbf{e}_1$ of the triangle $[\mathbf{w}] \cap c$, $\mathbf{e}_1$, $[\mathbf{v}] \cap c$. The meaning of this relation is that by restricting the point $\mathbf{z}_N$ to a smaller disk we can ensure a minimum angle enclosed by vectors $\mathbf{v}, \mathbf{w}$. It follows immediately that $\mathbf{z}_N = \mathbf{e}_1$

selects the unique orthogonal conjugate pair, i.e., the principal directions.

We use the word *transverse* to indicate that two lines do not coincide, and the degree of transversality is measured by the angle between these lines. A field of conjugate directions which are transverse everywhere is called a *TCD field.*

Singularities and indices of TCD fields. The discussion of singularities' properties is relevant for the basic question whether there exist TCD fields without them. The number and location of singularities of a TCD field is of course important if we later align a quad mesh with a TCD field: Singularities of the field will become combinatorial singularities of the mesh.

We start with the definition of *index*, which is illustrated by Figure 4. Assume that D is a simply connected domain in a surface, and that we have a continuous assignment of directions to each point of the boundary loop ∂D. When traversing ∂D in the positive sense, the direction has a total rotation angle ρ_D. The index ind_D is then defined by

$$\mathrm{ind}_D = \frac{1}{2\pi}\rho_D. \tag{2}$$

This index is always an integer multiple of $1/2$. A *cross field*, which is an assignment of an orthogonal pair of directions to each point, likewise has an index (see Figure 4) which is an integer multiple of $1/4$. For more details on indices of such fields the reader is referred to [RVLL08, RVAL09]. For a general field of transverse directions (such as a TCD field) we measure the index via its field of angle bisectors which is a cross field.

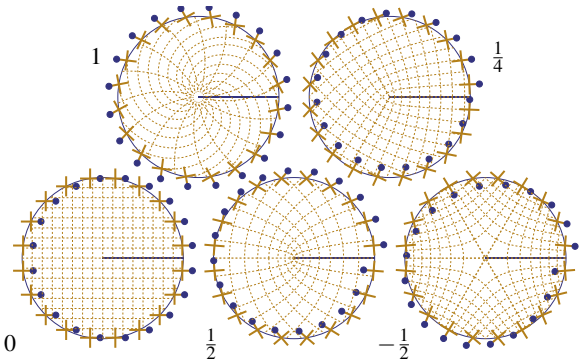

Figure 4: *The index of a field of transverse directions. The field is shown only along the boundary loop ∂D of the domain under consideration; inside D we show integral curves. We give the index ind_D which corresponds to a total rotation angle $2\pi\,\mathrm{ind}_D$.*

Def. 1. *The index $\mathrm{ind}_{\mathbf{p}}$ of a TCD field which is continuous around $\mathbf{p}$ (with the possible exception of $\mathbf{p}$ itself) is the index w.r.t. any small loop around that point. If it is nonzero, the field has a singularity at $\mathbf{p}$.*

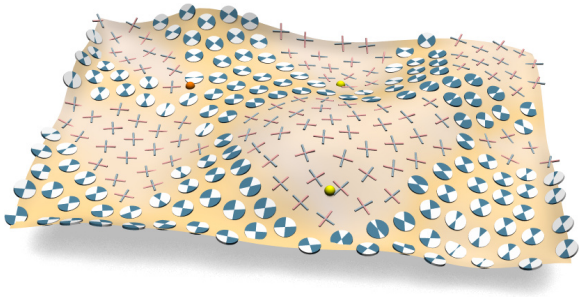

Figure 5: *'Flying carpet' surface, Louvre, Paris. We illustrate the degrees of freedom which the conjugate direction fields enjoy. In the negatively curved areas ($K<0$) we show the sectors where elements of such a TCD field are confined in. In areas where K approaches zero from below, one sector becomes thin and the corresponding member of a TCD field has not much freedom to move. In the positively curved areas ($K>0$), where TCD fields can rotate freely, we draw the major and minor principal direction as an example of a TCD field. Singularities of index $1/2$ and $-1/2$ are marked by small yellow resp. red balls.*

For an illustration of a TCD field and its indices, see Figure 5. It is well known that the index is well defined and additive when dissecting a domain into pieces:

$$\mathrm{ind}_D = \sum_{\mathbf{p}\in D} \mathrm{ind}_{\mathbf{p}}\,. \tag{3}$$

The field of principal curvature directions does not exhibit arbitrary indices. Only multiples of $1/2$ occur. This fact can be generalized:

Prop. 3. *A TCD field constructed by means of an auxiliary vector field $\mathbf{z}_N$ as described above always is the union of two separate direction fields; consequently all indices are integer multiples of $1/2$ (this especially applies to the principal direction field which corresponds to $\mathbf{z}_N = \mathrm{const.} = \mathbf{o}$).*

Proof: Discontinuities of the field occur for $\mathbf{z}_M = \mathbf{z}_N$, otherwise $N^{-1}M$ is no multiple of the identity matrix and has two different real eigenvalues. Since $\det N > 0$ we may w.l.o.g. assume that $N = \left(\begin{smallmatrix} 1 & n_{12} \\ n_{12} & n_{22} \end{smallmatrix}\right)$. This makes N uniquely and continuously dependent on the point $\mathbf{z}_N$. The directions which constitute the TCD field are defined by the eigenvectors of $N^{-1}M$ and can therefore be distinguished by belonging to the greater and the smaller eigenvalue. □

Remark: Our way of handling fields of conjugate directions prohibits the indices $\pm 1/4, \pm 3/4, \ldots$ and so not all possible such fields are treated. Restriction to integer multiples of $1/2$ however has the effect that the total number of singularities is reduced, which is a major goal anyway.

On the invariance of indices. The main results of this subsection are negative in character: they say that in many situations singularities are entailed by the geometry of the sur-

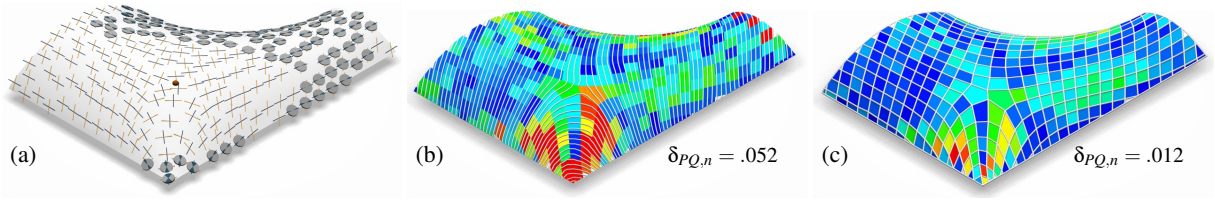

Figure 6: *Detail of the great court roof, British Museum, originally a triangle mesh design by Foster and Partners. (a) Design of a TCD field. In areas of nonpositive curvature the possible directions are indicated by sectors. (b) One half of integral curves of this field, found by means of our level set method. A quad mesh derived from these level sets is indicated by color coding the planarity measure $\delta_{PQ,n}$. (c) mesh already optimized for planarity of faces.*

face under consideration and cannot be removed. The first result concerns areas of negative Gaussian curvature, where the movement of conjugate directions is constrained by the asymptotic directions (see Figure 5):

Prop. 4. *In areas with $K < 0$ all TCD fields have the same indices and location of singularities.*

Proof: If $K < 0$, conjugate directions $[\mathbf{v}]$, $[\mathbf{w}]$ are separated by the self-conjugate *asymptotic* directions $[\mathbf{a}]$, $[\mathbf{b}]$ (see Figure 2). Thus both $[\mathbf{v}]$, $[\mathbf{w}]$ can possibly move only in their own respective sector which is bounded by $[\mathbf{a}]$, $[\mathbf{b}]$. It follows that for a loop inside the $K < 0$ area, the total rotation angle of the field equals the total rotation angle of either $[\mathbf{a}]$ or $[\mathbf{b}]$, which is independent of the field. □

Figure 2, right, and Figure 3, right imply that along the parabolic curves (which separate the positively curved areas from the negatively curve ones) there is even less freedom:

Prop. 5. *In all surface points where $K = 0$, one of two conjugate directions always equals the principal curvature direction corresponding to zero curvature.*

Prop. 6. *In areas where $K \leq 0$, all indices of TCD fields are integer multiples of $1/2$ (regardless of their way of construction via $\mathbf{z}_N$ or otherwise).*

Proof: Prop. 4 and Prop. 5 say that the indices of a TCD field coincide with the indices of the principal direction field. The statement now follows from Prop. 3. □

TCD fields have more freedom in the positively curved areas of a surface. It turns out, however, that the number of singularities present is in some cases bounded from below:

Prop. 7. *Assume that D is a connected component of the $K > 0$ area in the surface which does not touch the surface's boundary and which has no flat point with $\kappa_1 = \kappa_2 = 0$ on ∂D. Then the total index ind_D is the same for all TCD fields. Only if $\mathrm{ind}_D = 0$ we may have a TCD field without singularities in D.*

Proof: On ∂D one principal curvature vanishes. The corresponding principal curvature direction is part of any pair of conjugate directions (see Figure 2, right). It follows that the total rotation angle of the TCD field equals the angle of the principal field. □

Remark: In our implementation all elements of a TCD field $(\mathbf{z}_M, \mathbf{z}_N, \dots)$ are associated with the faces of a mesh. Deciding whether a face represents a singularity requires to compute the total rotation angle of the field along a path around f, for which we use a 1-neighbourhood. Because of Prop. 3 each of the two directions $\mathbf{v}$, $\mathbf{w}$ in the TCD field must result in the same rotation angle. In fact we use a locally defined auxiliary vector bisecting $\mathbf{v}$, $\mathbf{w}$ for computing the rotation angle.

4. Design and optimization of TCD fields.

The guiding principle behind the optimization of a field of conjugate directions is that we represent this field by a simple object, namely the vector field $\mathbf{z}_N$. We express all desired properties of the TCD field in terms of the vector field. For the handling of vector fields we employ 1-forms as proposed by [FSDH07].

We set up a target functional whose minimization is to ensure desirable properties of a TCD field, such as smoothness, transversality, absence of singularities in general, singularities in prescribed places, and prescribed values of the TCD field in some places.

Setup of global optimization. *Smoothness* of the TCD field is measured by smallness of the Dirichlet energy of the vector field $\mathbf{z}_N$ which is is computed according to [FSDH07]. We take that energy as the base of a target functional for optimization. For *transversality* we penalize small angles between $\mathbf{v}$ and $\mathbf{w}$ which according to (1) can be done by adding

$$\lambda_{\mathrm{trans}} \sum_{f \in F} \phi(\|\mathbf{z}_{N,f} - \mathbf{e}_{1,f}\|), \quad \phi(t) = t^4$$

to the target functional, where ϕ can be any function which grows quickly as we approach 1. If in a face f we wish to place a singularity we can simply enforce $\mathbf{z}_{M,f} = \mathbf{z}_{N,f}$. In case that in a certain area we wish to discourage formation of singularities, we prevent $\mathbf{z}_M = \mathbf{z}_N$ by adding to the target

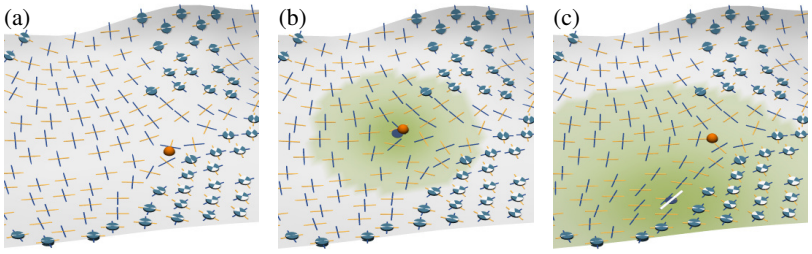

Figure 7: *Local corrections to the TCD field shown in (a). (b) We move a singularity to a new desired position (shown by the location of the red ball). (c) Local deformation of field to achieve desired direction of 1 conjugate direction (white). The green area indicates the extent of the sub-mesh involved.*

functional the term

$$\sum_{f\in F} \lambda_{reg,f}\psi(\|\mathbf{z}_{N,f}-\mathbf{z}_{M,f}\|), \quad \psi(t)=\begin{cases}(t-1)^4, & \text{if } t<1,\\ 0 & \text{else.}\end{cases}$$

Prop. 2 shows us how to encode, in terms of $\mathbf{z}_{N,f}$, the condition that the conjugate vectors $\mathbf{v}_f, \mathbf{w}_f$ assume prescribed values: We must penalize deviation of $\mathbf{z}_{N,f}$ from the straight line $= ([\mathbf{v}_f]\cap c_f)\vee([\mathbf{w}_f]\cap c_f)$ (using the notation of the figure close to Prop. 2). If required, we add the square of this distance to the target functional.

Implementation. Optimization is done by a nonlinear conjugate gradient method [NW99]. We used the Fletcher-Reeves way of updating descent directions, but this choice is not critical. The necessary first derivatives are computed numerically. To initialize optimization we first reset to zero all coefficients of the 1-form representing the vector field $\mathbf{z}_N$, and next optimize for $\sum_{f\in F}\|\mathbf{z}_{N,f}-\mathbf{e}_{1,f}\|^2 \to \min$. This achieves $\mathbf{z}_N \approx \mathbf{e}_1$ (the corresponding TCD field is as close to principal as possible). The vector field $\mathbf{z}_N$ constructed in this way is used as a starting point for minimizing the target functional assembled above.

Local correction: Moving singularities. The moving of a singularity of a TCD field from a current position f_0 to a nearby new position f_1 can be efficiently and interactively performed without resorting to global optimization. With the smooth function

$$\sigma(t)=\begin{cases}1 & \text{if } t\in[0,r_1],\\ \frac{1}{2}+\frac{1}{2}\cos(\pi\frac{t-r_1}{r_2-r_1}) & \text{if } t\in[r_1,r_2],\\ 0 & \text{if } t\in[r_2,\infty),\end{cases} \tag{4}$$

we perform a smooth correction to the vector field $\mathbf{z}_N$ and replace it by

$$\mathbf{z}^{\text{new}}_{N,f}=\mathbf{z}_{N,f}-\beta_f(\mathbf{z}_{N,f_1}-\mathbf{z}_{M,f_1}), \quad \text{where}$$
$$\beta_f=\sigma(\text{dist}(f,f_0)).$$

This computation is two-dimensional; in each face we use the appropriate local coordinate system. Assuming that the face f_1 is close to f_0 so that $\beta_{f_1}=1$, we then have

$$\mathbf{z}^{\text{new}}_{N,f_1}=\mathbf{z}_{N,f_1}-(\mathbf{z}_{N,f_1}-\mathbf{z}_{M,f_1})=\mathbf{z}_{M,f_1}.$$

Because we have now achieved $\mathbf{z}_{M,f_1}=\mathbf{z}^{\text{new}}_{N,f_1}$, the singularity of the TCD field has moved to its desired new location f_1. Figure 7a,b illustrates this procedure which is intended to be worked with interactively.

Local optimization. We should point out that optimization can be localized by simply applying it to part of the mesh, using functions like (4) to blend a locally modified vector field with the original one (see Figure 7c for an example).

5. Meshing via level sets.

In order to set up a quad-dominant mesh whose edges are aligned with a TCD field, it is very convenient to first find functions defined on the surface Φ, whose level sets are aligned with that TCD field. It is also possible to incorporate our wish for an even spacing of vertices into this level set formulation. It is not difficult to set up this level set method in a simply connected area where the TCD field is regular: at singularities some additional considerations will be necessary.

A quad mesh whose edges are aligned with the level sets of functions g,h which are defined in the given surface Φ is easily found by choosing values w_i for the function g and w'_j for the function h, and placing a vertex $\mathbf{v}_{ij}$ such that

$$g(\mathbf{v}_{ij})=w_i, \quad h(\mathbf{v}_{ij})=w'_j.$$

It is natural to choose $w_i=i\cdot\Delta w$, and $w'_j=j\cdot\Delta w'$ (i.e., equally spaced samples). It remains to choose the functions g and h.

Optimization Setup. If we are given a conjugate direction field without singularities, then we can globally represent this field using two nonzero vector fields $\mathbf{v}$, $\mathbf{w}$. The unknown functions g,h are considered to be piecewise-linear on the mesh, so they are determined by their values in the vertices. Their gradient vector fields are piecewise-constant. In the following we assume that Φ is simply connected. We optimize the values of g such that

$$\sum\nolimits_{f\in F}\text{area}(f)\|R_f\mathbf{v}_f-(\nabla g)_f\|^2\to\min, \tag{5}$$

where R_f is the rotation about 90 degrees in the face f. A similar formula applies to the function h and the vector field $\mathbf{w}$. An exact solution of (5) which then has $\nabla g = R\mathbf{v}$ exists if and only if the vector fields $\{R_f\mathbf{v}_f\}_{f\in F}$ and $\{R_f\mathbf{w}_f\}_{f\in F}$ are integrable. Only in this case will the level sets of g resp. h

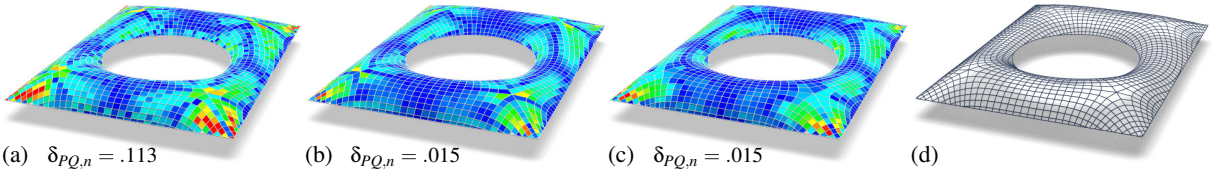

(a) $\delta_{PQ,n} = .113$ (b) $\delta_{PQ,n} = .015$ (c) $\delta_{PQ,n} = .015$ (d)

Figure 8: *Great court roof, British Museum, originally a triangle mesh design by Foster and Partners. The quad mesh in (a) is found by our level set method from the principal directions, which exhibit four singularities, each of index $-1/2$. Color coding shows the degree of planarity of each face. (b) optimization for the PQ property, and (c) optimization for the conical property [LPW*06]. (d) shows the mesh of (c). The fact that all these meshes are pretty much the same proves that initializing the first mesh from principal directions is near-optimal.*

align exactly with the vector fields $\mathbf{v}$, $\mathbf{w}$. Note that the *length* of vectors $\mathbf{v}_f$, $\mathbf{w}_f$ influences the integrability. We use these lengths on purpose, because they allow us to control the distance of successive level sets $[g = w_i]$ and $[g = w_{i+1}]$: this distance approximately equals

$$(w_{i+1} - w_i)/\|\nabla g\|.$$

Thus we can incorporate the very important property of *equal spacing* of vertices in our method by setting the vectors $\mathbf{v}_f$, $\mathbf{w}_f$ to unit length. The minimization problem (5) only involves the gradient of g and thus determines g only up to a constant: we have to fix one value of g.

Implementation. Since Equ. (5) is quadratic, optimization amounts to solving a linear equation system for the values of g at the vertices. Below we encounter further functionals to minimize which are no longer quadratic (Equations (8), (7) and (6)). We employ Gauss-Newton for their optimization. All required first order derivatives are computed exactly. The linear systems to be solved in each iteration step are sparse, because all contributions to the functionals are local. We can therefore employ sparse Cholesky factorization using CHOLMOD [CDHR08].

Remark: The fact that the length of vectors enters (5) may lead to a bad alignment of level sets with the given conjugate directions. If we do not care about equal spacing at all, we replace (5) by

$$\sum_{f \in F} \frac{\text{area}(f)}{\text{area}(\Phi)} \left\langle \mathbf{v}_f, \frac{\nabla g_f}{\|\nabla g_f\|} \right\rangle^2 \to \min. \qquad (6)$$

We minimize this nonlinear functional by a Gauss-Newton Method, using a solution of (5) for initialization.

Quad meshes from a TCD field with singularities. In case the given TCD field exhibits singularities, we can not globally represent it by two nonzero vector fields. This can be remedied conceptually by lifting the TCD field to a suitable branched covering Φ' of the given surface Φ, as described in [KNP07]. In our implementation this amounts to introducing cuts which run from singularities to the surface boundary. Functions g, h whose level sets are to be aligned with vector fields now exhibit jumps across cuts; and these jumps must be defined in a way which makes level sets continuous.

This is implemented as follows: A cut is represented as an edge polyline $\mathscr{C}$ in the triangle mesh. Consider first a cut which emanates from a singularity of index $\pm 1/2$. Using the notation $g^{\text{left}}(v)$ and $g^{\text{right}}(v)$ for the two function values of a vertex $v \in \mathscr{C}$, we must have

$$g^{\text{left}}(v) = \alpha_{\mathscr{C}} - g^{\text{right}}(v), \quad \text{where} \quad \alpha_{\mathscr{C}} = i_{\mathscr{C}} \cdot \Delta w$$

for all vertices $v \in \mathscr{C}$ and some integer $i_{\mathscr{C}}$. For the source vertex v_0, the two function values coincide, which implies

$$g(v_0) = \alpha_{\mathscr{C}}/2.$$

Further cuts may be necessary to make the surface simply connected (e.g. for the model of Figure 8). Here we have

$$g^{\text{left}}(v) = \alpha_{\mathscr{C}} + g^{\text{right}}(v), \quad \text{where} \quad \alpha_{\mathscr{C}} = i_{\mathscr{C}} \cdot \Delta w$$

for all vertices $v \in \mathscr{C}$ and some integer $i_{\mathscr{C}}$. A similar equation applies to the function h and the step size $\Delta w'$. These conditions guarantee continuity of level sets across cuts.

Remark: Higher order singularities of index $\pm k/2$ can be seen as the limit of k singularities of index $\pm 1/2$, with k cuts emanating from the singularity which run parallel. Here we have the relation $g^{\text{left}}(v) = \alpha_{\mathscr{C}} + (-1)^k g^{\text{right}}(v)$, but the singular source vertex of the cut is, in theory, associated with k function values. Our examples have only $k = 1$.

The optimization problems (5) and (6) are modified as follows: Vertices which take part in a cut have two function values related by the linear jump conditions given above. We introduce the constants $\alpha_{\mathscr{C}}$ as new variables and first optimize without restricting the values which the constants $\alpha_{\mathscr{C}}$ can assume. In a subsequent step we are rounding the constants $\alpha_{\mathscr{C}}$ which belong to cuts emanating from a singularity to the nearest *even integer* multiple of the appropriate step sizes Δw or $\Delta w'$ in order to achieve an all-quad meshing where level sets pass through the singularities. If we round to the nearest *odd integer* multiple we get only vertex 4 vertices but extraordinary faces instead. We likewise round the constants $\alpha_{\mathscr{C}}$ associated with the remaining cuts. Now optimization is repeated, where all values achieved by rounding are kept constant.

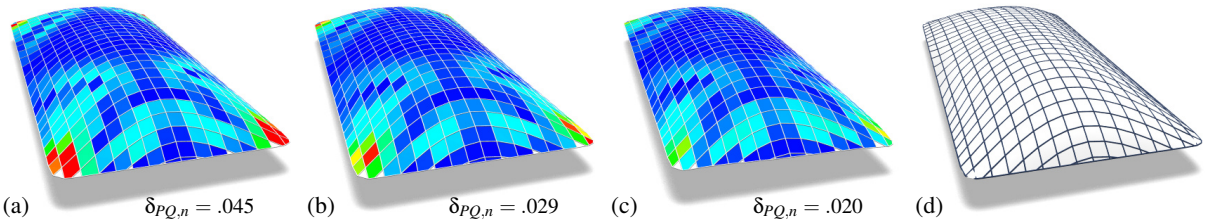

Figure 9: *A design study for the courtyard roof of Neumünster monastery, originally a triangle mesh design by RFR which follows a surface exhibiting a slight tangent discontinuity. One family of mesh polylines consists of planar sections of the surface. We show optimization towards planarity with different side conditions. (a) before optimization. (b) optimization taking into account deviation from the reference surface and deviation of boundaries (c) optimizing without regard to boundary. (d) the mesh of (b) overlaid on the original surface.*

Quad meshes without TCD fields. In special cases it is possible to prescribe one family of level sets explicitly. For instance, Figure 9 shows a shape Φ covering a rectangular courtyard where we prescribe that the level sets of g are planar sections of Φ. We could also formulate this condition in terms of TCD fields (by requiring that one of the two conjugate directions is parallel to a fixed plane), but here we can avoid TCD fields entirely. We consider the function g as given, and compute the function h as a minimizer of the following functional, such that the level sets of g, h have conjugate tangents:

$$\sum_{f\in F} \frac{\text{area}(f)}{\text{area}(\Phi)} \left(\left(R\frac{\nabla g}{\|\nabla g\|} \right)_f^T \cdot M_f \cdot \left(R\frac{\nabla h}{\|\nabla h\|} \right)_f \right)^2. \qquad (7)$$

Here we have abused notation and assumed that all involved vectors are represented in the local coordinate frame associated with the face f which was discussed in Section 2. We augment this target functional by $\lambda_{\text{even}}\,\text{area}(\Phi)\sum_{f\in F}\|(\nabla h)_f\|^2$ in order to achieve equal spacing of level sets. The latter functional is quadratic and minimizers are easily found if the values of h at two selected vertices are fixed. We use such a minimizer to initialize minimization of (7).

6. Optimization of quad meshes for planarity.

The quad-dominant meshes which are the result of the procedures described above are subject to further optimization in order to make their faces planar. Since we took the available curvature information into account when we created those meshes in the first place, we are already close to exact planarity. For this optimization we set up the functional

$$f_{PQ} + \lambda_{fair} f_{fair} + \lambda_{prox} f_{prox} + \lambda^{\partial}_{prox} f^{\partial}_{prox}. \qquad (8)$$

The definition of the single contributions to (8) uses the notation diag_f for the distance between diagonals of the quadrilateral face f, V_∂ for the set of boundary vertices, π and $\widetilde{\pi}$ for the closest-point projections onto Φ and $\partial\Phi$, respectively; τ_p for the tangent plane in p, and finally T_p for the tangent of the boundary curve $\partial\Phi$ in the point p. We define as follows:

$$f_{PQ} = \sum_{f\in F} (\text{diag}_f)^2, \quad f_{prox} = \sum_{v\in V\setminus V_\partial} \|v - \tau_{\pi(v)}\|^2,$$
$$f^{\partial}_{prox} = \sum_{v\in V_\partial} \|v - T_{\widetilde{\pi}(v)}\|^2.$$

Further, we measure fairness by comparing second differences "Δ^2_{uvw}" of every triple u, v, w of consecutive vertices with the respective original value before optimization:

$$f_{fair} = \sum_{\text{triples } (uvw)} \|\Delta^2_{uvw} - \Delta^{2,\text{orig}}_{uvw}\|^2.$$

Details of optimization of (8) are summarized by Figure 10. The computation time refers to a dual core 2.4GHz MacBook Pro. The quality of results is indicated by the proximity measures

$$\delta_{prox} = \max_{v\in V\setminus V_\partial} \|v - \tau_{\pi(v)}\|, \quad \delta^{\partial}_{prox} = \max_{v\in V_\partial} \|v - T_{\widetilde{\pi}(v)}\|,$$

the planarity measure $\delta_{PQ} = \max_{f\in F} \text{diag}_f$, and the normalized planarity measure $\delta_{PQ,n}$, which is the maximum of diag_f divided by the mean length of diagonals in the face f.

Model	# Var.	# Iter.	sec	δ_{PQ}	$\delta_{PQ,n}$	δ_{prox}	δ^{∂}_{prox}
Figure 1	5373	32	18.1	.0003	.011	.0003	.0003
Figure 6	1302	16	3.0	.0007	.012	.0003	.0008
Figure 8	3462	39	15.1	.0003	.01	.0003	.0001
Figure 9	1539	13	2.7	.0013	.029	.0004	.0006

Figure 10: *Details of PQ optimization. Bounding box diameter of all objects equals* 1.

7. Discussion and limitations.

Testing our methods on datasets which originate in real-world freeform architectural designs such as demonstrated by the figures in this paper showed satisfactory results. One can see that optimization towards planarity reduces the normalized planarity measure $\delta_{PQ,n}$ only by a factor five or so,

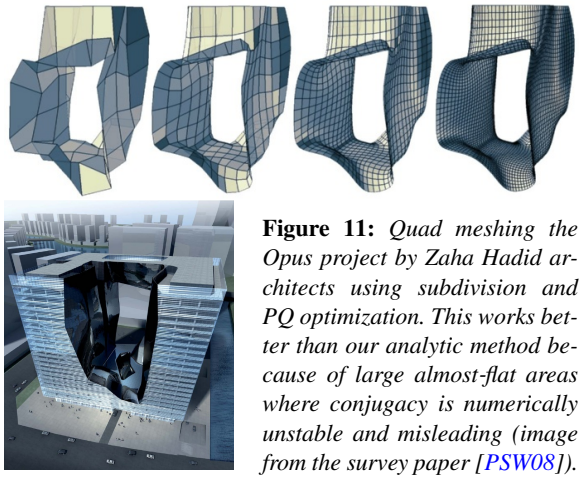

Figure 11: *Quad meshing the Opus project by Zaha Hadid architects using subdivision and PQ optimization. This works better than our analytic method because of large almost-flat areas where conjugacy is numerically unstable and misleading (image from the survey paper [PSW08]).*

confirming that alignment of mesh and TCD field is a good initialization for the final round of optimization.

We should mention that the many changes between positive and negative curvature in shapes like Figure 5 forces PQ meshes to basically follow the principal curvature lines with almost no degrees of freedom left to achieve even spacing of vertices. This is not a a limitation of the method, but a limitation of the design. In fact the analytic method of the present paper is useful to detect such features of a design.

There are instances where our method is not optimal: for an example and a reason see Figure 11. An even simpler example where our method fails completely is a cube which is made smooth by rounding edges and corners. This is due to the lack of principal directions in most of this surface.

Perhaps the most severe restriction of our method is that we handle only TCD fields whose indices are integer multiples of $1/2$. In the positively curved areas of a surface there exist other TCD fields whose indices are only integer multiples of $1/4$. Especially when meshing convex corners, fields with index $1/4$ are natural, and we again identify the rounded cube as a failure case of our method.

However, the paper's method is not meant for such situations. It is meant to deal with the so-called "second order" features in surfaces which appear smooth and rounded to the eye, without the eye being able to detect the obstructions to meshing with planar quads. With the absence of first order features such as present in a rounded cube, the need for singularities with index $1/4$ disappears and the goal of as few singularities as possible dominates.

It is possible to combine different methods in a pragmatic way, for instance by using the paper's method to find a mesh for those parts of a surface where principal curvature information exists, and using this information to redo meshing by the subdivision approach – knowing already where to put extraordinary vertices, and knowing that in the flat and near-spherical parts, mesh combinatorics is irrelevant for planarity.

Acknowledgments

The research leading to these results has received funding from the European Community's Seventh Framework Programme under grant agreement No. 230520 ("ARC"), and grant No. 813391 of the Austrian research promotion agency (FFG). We want to express our thanks to Waagner Biro Stahlbau (Vienna) and RFR (Paris) for datasets of architectural designs.

References

[BZK09] BOMMES D., ZIMMER H., KOBBELT L.: Mixed-integer quadrangulation. *ACM Trans. Graphics 28*, 3 (2009), #77, 1–10. 2

[CAD04] COHEN-STEINER D., ALLIEZ P., DESBRUN M.: Variational shape approximation. *ACM Trans. Graphics 23*, 3 (2004), 905–914. 2

[CDHR08] CHEN Y., DAVIS T. A., HAGER W. W., RAJAMANICKAM S.: Algorithm 887: CHOLMOD, supernodal sparse Cholesky factorization and update/downdate. *ACM Trans. Math. Softw. 35*, 3 (2008), #22, 1–14. 7

[Cox92] COXETER H. S. M.: *The real projective plane*, 3rd ed. Springer, 1992. 3

[CP03] CAZALS F., POUGET M.: Estimating differential quantities using polynomial fitting of osculating jets. In *Symp. Geometry processing* (2003), Kobbelt L., Schröder P., Hoppe H., (Eds.), Eurographics, pp. 177–178. 3

[dC76] DO CARMO M.: *Differential Geometry of Curves and Surfaces*. Prentice-Hall, 1976. 2

[FSDH07] FISHER M., SCHRÖDER P., DESBRUN M., HOPPE H.: Design of tangent vector fields. *ACM Trans. Graphics 26*, 3 (2007), #56, 1–10. 3, 5

[KNP07] KÄLBERER F., NIESER M., POLTHIER K.: QuadCover – surface parameterization using branched coverings. *Comput. Graph. Forum 26*, 3 (2007), 375–384. 7

[LPW*06] LIU Y., POTTMANN H., WALLNER J., YANG Y.-L., WANG W.: Geometric modeling with conical meshes and developable surfaces. *ACM Trans. Graphics 25*, 3 (2006), 681–689. 1, 2, 7

[NW99] NOCEDAL J., WRIGHT S. J.: *Numerical Optimization*. Springer, 1999. 6

[PLW*07] POTTMANN H., LIU Y., WALLNER J., BOBENKO A., WANG W.: Geometry of multi-layer freeform structures for architecture. *ACM Trans. Graphics 26*, 3 (2007), #65, 1–11. 1

[PSB*08] POTTMANN H., SCHIFTNER A., BO P., SCHMIEDHOFER H., WANG W., BALDASSINI N., WALLNER J.: Freeform surfaces from single curved panels. *ACM Trans. Graphics 27*, 3 (2008), #76, 1–10. 2

[PSW08] POTTMANN H., SCHIFTNER A., WALLNER J.: Geometry of architectural freeform structures. *Int. Math. Nachr. 209* (2008), 15–28. 9

[RVAL09] RAY N., VALLET B., ALONSO L., LÉVY B.: Geometry-aware direction field processing. *ACM Trans. Graphics 28*, 1 (2009), #1, 1–13. 2, 4

[RVLL08] RAY N., VALLET B., LI W. C., LÉVY B.: N-symmetry direction field design. *ACM Trans. Graphics 27*, 2 (2008), #10, 1–13. 2, 4