Statics-Sensitive Layout of Planar Quadrilateral Meshes

Alexander Schiftner

TU Wien, Evolute

Jonathan Balzer

KAUST

Abstract. *Rationalization of architectural freeform structures using planar quadrilateral (PQ) meshes has received rising interest in recent years, facilitated mainly by the introduction of algorithms which are capable of generating such. These algorithms involve an optimization which is, up to now, motivated purely geometrically and accounts for aspects of feasibility, visual appearance, and approximation of the architectural design. Practitioners would wish to add stiffness to the objectives of the layout process. This paper presents a simple but effective statics-aware initialization procedure for the layout of PQ meshes approximating a given freeform surface. We focus on the class of surface structures with membrane-like load bearing behavior, quite regularly encountered in architecture. By compliance analysis of two representative examples, we demonstrate that this specific type of initialization has indeed favorable impact on the mechanical properties of the final PQ mesh.*

1 Introduction

1.1 Motivation

A common approach for rationalizing architectural freeform surfaces using a collection of flat panels is by means of triangle meshes, like the example in Figure 1. Recent research opened up the possibility of approximating architectural freeform surfaces by PQ meshes that exhibit sufficient quality for the architectural application in terms of visual appearance and practical size requirements. Besides the observation that PQ meshes admit torsion-free nodes and may therefore be realized more light-weight, little has been said so far about their load-bearing behavior. The aim of this contribution is to close this gap.

1.2 Outline of the Proposed Method

It is well known that the geometric constraints on the prospective layouts of a PQ mesh are substantially more profound than those of a triangular one. We thus believe that the yet complex system of different objectives could well cloud if not entirely

Figure 1: Roof over the Great Courtyard of the British Museum, London, in classical realization by triangular steel/glass panels (Source: http://carpelondinium.wordpress.com). We present a rationalization by a planar quadrilateral mesh with optimized statics.

compensate the effect of an additional stiffness-enforcing component. This is even more true as inconsistencies among the objectives must be expected. Luckily, the present work founds on an algorithm which operates in the spirit of classical perturbation approaches insofar that, roughly speaking, the non-planar quad mesh is shifted towards a proximate local minimizer of improved planarity. It is therefore reasonable to assume that the planarization result will, to some extent, inherit the mechanical properties of the initial mesh. This immediately implies the following two-step strategy:

1. Instead of optimizing the network of rods as part of the remeshing step explained next, we treat the given reference geometry as a continuum. Although the "discrete" shape of the structure (e.g. the sizing or density of rods) is discarded early in process, this presents the only possibility to perform the desired decoupling of statical optimization and mesh layout. The outcome of this "homogenization" is a so-called *shell*, which is, loosely speaking, a regular surface endowed with a selection of mechanical properties. We compute its strain state for the case that only gravitational forces induced by self-weight act upon the structure. After subsequent principal component analysis, a vector field of maximal loading directions can be passed to the meshing module.

2. We lay out a (quasi)-planar quad mesh which is aligned with the precomputed guidance vector field. Regular quad meshes approximating a freeform surface are characterized by two families of edge polygons which may be considered as discrete versions of smooth curves on the underlying

freeform surface. Requiring the quads to be planar poses a constraint on the two families of edge polygons respectively curves. They have to resemble a so-called *conjugate curve network*, see [Liu et al. 2006] for a definition. We use the guidance vector field resulting from the mechanical preprocessing phase as an input for computing a function defined on the given freeform surface, whose level set curves are aligned closely with the vector field. Furthermore, we compute a second function whose level set curves are conjugate to the level set curves of the first function. Having computed such functions, we may extract a quad-dominant mesh (a mesh consisting mainly of quads) from them, whose faces are not exactly planar but close to. Subsequent optimization of these meshes towards planarity of faces can be expected to be successful as we illustrate by examples. Furthermore, the resulting PQ meshes are likely to exhibit improved load bearing behavior compared to other PQ meshes approximating the same freeform surface.

1.3 Related Work

Mechanics of Thin-Shell Structures. Shell structures and their accurate as well as computationally tractable modeling remain an active area of research. The relevant sources are far too numerous to provide an exhaustive list here. Instead, we point the reader to the pioneering works by Koiter [Koiter 1970], Naghdi [Naghdi 1972], Reisner [Reissner 1947], Kirchhoff, and Love [Love 1888] as initial references. Recently, there has been an increased interest in thin shells in the graphics and computational geometry communities, mainly motivated by the aim to create physically sound deformation models for surface meshes, see e.g. [Grinspun et al. 2003] and in particular the survey [Botsch and Sorkine 2008]. The paper [Winslow et al. 2008] is probably the one most closely related to ours. The authors point out the connection between macroscopic grids in architectural geometry and – at a microscopic scale – fibres in a reinforced polymer composite. As the present contribution, they rely on a homogenization approach. But although, for static computations, they treat quadrilateral cells of the target grid as plates, i.e. planar shells, they do not seem to enforce a planarity condition in their multi-criterial grid optimization framework.

Planar Quadrilateral Meshes. Early applications of PQ meshes for architectural purposes can be seen in projects by Schlaich-Bergermann. These PQ meshes are discrete versions of translational, rotational, and a few further types of surfaces, refer to [Glymph et al. 2002]. More recent research initiated by [Liu et al. 2006] deals with the possibility of approximating general freeform surfaces by PQ meshes. The authors introduce an optimization method which perturbs the vertex positions of an initial quad mesh such that the faces become planar (so-called *planarization*). In order to generate a suitable initial quad mesh, we use a surface parametrization or level set method. Numerous papers have been published on this topic in previous years. The method we use is described in detail in [Zadravec et al. 2010] and closely related to [Kälberer et al. 2007].

2 Statics-Sensitive Preprocessing

We will now explain our simple mechanical model of the geometry to be remeshed. Section 2.2 discusses the numerical solution of the resulting partial differential equation (PDE), Section 2.2 how it interacts with the subsequent meshing step. All the while, it is not our intent to estimate the true stress and strain state of the surface under consideration as accurately as possible but rather provide a qualitative answer to the question how the multiply constrained meshing problem can be influenced positively in a statical sense.

2.1 Thin-Shell Model

A *shell* is a three-dimensional mechanical structure, but it is slender in the sense that its thickness is small compared to its lateral dimensions. It may be seen as a stack of two-dimensional layers, all distinguished by their distance to the so-called *mid-surface*. We associate with the mid-surface of a shell the given geometry that we want to approximate by a PQ mesh. The underlying theory is well-established and can be reviewed in the survey [Botsch and Sorkine 2008] among numerous other sources. In short: The constitutive equation of shell deformation derives from a simple conservation law. In static equilibrium, the internal elastic energy E_{el} equals the mechanical work W applied to the body by external forces $\boldsymbol{f}$. The former dissolves into two parts: one which is induced by shearing or stretching the structure in tangential directions (the so-called *membrane energy*) and the other one which is due to bending. In what ratio the two deformation modes appear in a structure depends mainly on its bearing, i.e. the boundary conditions, and sometimes the prevailing load case, see Figure 2. A linearized deformation model, i.e. a vector field $\boldsymbol{v}$ of displacements between the original shell and its deformation, is found by solving the Euler equation of the corresponding energy residual functional $E_{\mathrm{el}} - W$:

$$-\lambda_{\mathrm{m}}\Delta\boldsymbol{v} + \lambda_{\mathrm{b}}\Delta^2\boldsymbol{v} = \boldsymbol{f}. \tag{1}$$

The proportionality constants λ_{m} and λ_{b} are material parameters describing deformation resistance or stiffness with respect to membrane actions respectively bend-

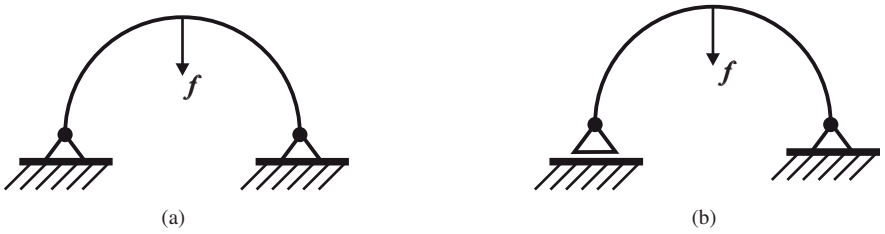

Figure 2: (a) A pre-stressed bearing fosters tangential force flow i.e. bending-stiff membrane behavior. (b) Introduction of a floating bearing significantly reduces the resistance to bending.

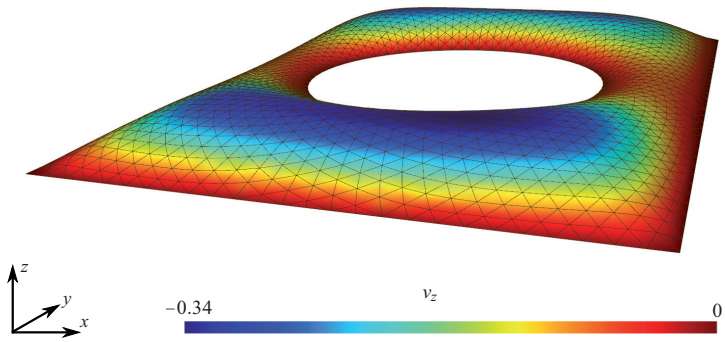

Figure 3: For mechanical preprocessing, the reference geometry is modeled as a membrane shell. The z-component v_z of $\boldsymbol{v}$ owing to self-weight is obtained by solving, in this special example, a Dirichlet problem (fixed boundary). The color encodes its distribution over the undeformed surface.

ing. The Laplace and bi-Laplace operator Δ respectively Δ^2 each represent one of the two deformation modes.

2.2 Implementation Issues

Estimating the Strain State For a discretization of the reference geometry a manifold mesh is used, i.e. a triangulation without self-intersections, hanging edges, t-junctions, or isolated vertices. This choice is very favorable with regard to the finite-elements (FE) formulation of (1): Locking effects i.e. insufficient convergence rates, typical complications of thin shell numerics, are not to be expected for this particular type of element, cf. [Hughes 2000]. Galerkin projection on its nodal basis situates Equation (1) in the finite-dimensional space of functions whose restriction to each individual triangle is linear. The outcome is a sparse linear system of algebraic equations, which can be solved rather cheaply with the help of standard software, e.g. the CHOLMOD library, featuring supernodal column ordering and sparse Cholesky decomposition [Davis and Hager 2009].

In theory, our method includes bending-dominated deformations, which cause, however, significant numerical complications arising at discretization of the bi-Laplacian Δ^2 in (1). One could be tempted to simply square an approximation of the Laplacian Δ. This practice is somewhat questionable as fourth-order PDEs necessitate designing a set of C^1-continuous shape functions, which the piecewise linear basis is clearly not. For a remedy, one would either have to increase the element function order, i.e. perform a p-refinement, or rely on different discretizations such as e.g. subdivision surfaces, cf. [Cirak et al. 2000], alternatively B-spline patches like in isogeometric analysis [Hughes et al. 2005]. Realization of any aforementioned strategy is well beyond the scope of the present work. Consequently, we omit the fourth-order term for the time being, also see the concluding remarks in

Section 5.

The current implementation features the possibility to assign either a Dirichlet or Neumann condition to every boundary node. As depicted in Figure 4, the former restricts translational, the latter rotational degrees of freedom. Although in practice, a building is subject to a variety of different forces, we confine ourselves to the case of dead loading[1]. Inclusion of more general scenarios is straightforward. The force field $\boldsymbol{f}$ is concentrated on the vertex set of the mesh by means of a simple mass lumping. Observe that, since in this case the x- and y-components of $\boldsymbol{f}$ are zero and, by the maximum principle, also those of $\boldsymbol{v}$, the system (1) reduces to a single scalar equation in the z-component of the latter.

(a) (b)

Figure 4: Available boundary conditions: (a) Setting $\boldsymbol{v} = \mathbf{0}$ on the boundary amounts to a fixed but limp bearing. (b) Enforcing $D\boldsymbol{v}\hat{\boldsymbol{o}} = \mathbf{0}$, where $\hat{\boldsymbol{o}}$ is the outer normal of Γ, suppresses rotations around the boundary curve but does not necessarily limit translational degrees of freedom.

Principal Strains Once $\boldsymbol{v}$ has been recovered from (1), we evaluate the so-called *Green Lagrange tensor*

$$\boldsymbol{\varepsilon}_{\mathrm{m}} := \frac{1}{2}(DT^{\top}DT - \mathbf{P}) \tag{2}$$

in a face of the mesh as follows: First, the displacement $\boldsymbol{v}$ is converted into the vector-valued mapping T between deformed and original configuration of the shell. In the discrete setting, a scalar function in the piecewise linear FE space basis is uniquely defined by its values on the node set. It is thus straightforward to compute its gradient, which is constant on each single face, from the values on the three incident vertices. Concatenating the face-based gradients of all components of T then approximates the Jacobian DT. All operations are performed with respect to a fixed Eulerian world coordinate system. Hence, the Green-Lagrange tensor $\boldsymbol{\varepsilon}_{\mathrm{m}}$ acquires the shape 3×3 and a nontrivial kernel, spanned by the current unitary face normal $\hat{\boldsymbol{n}}$. For this particular choice of basis, the projection matrix $\mathbf{P} := \mathbf{I} - \hat{\boldsymbol{n}}\hat{\boldsymbol{n}}^{\top}$ restricted to the tangent space corresponds to the identity there. After putting intermediate results together, it is illustrative to verify that, given a direction $\boldsymbol{t} \in \mathbb{R}^3$, intra-surface length discrepancy, i.e. membrane strain, is quantified via $\boldsymbol{t}^{\top}\boldsymbol{\varepsilon}_{\mathrm{m}}\boldsymbol{t}$.

By eigenvalue decomposition of $\boldsymbol{\varepsilon}_{\mathrm{m}}$, we get the orthonormal and symmetric frame field $(\boldsymbol{e}_{\max}, \boldsymbol{e}_{\min})$. It is ordered and labeled by magnitude of the associated

[1] One usually differentiates between dead and live loads. The former is attributed to pure self-weight. The latter include acceleration forces due to non-stationary winds, earthquakes, or similar as well as time-dependant mass surpluses in the form of snow, rain water, traffic if the structure is e.g. a bridge.

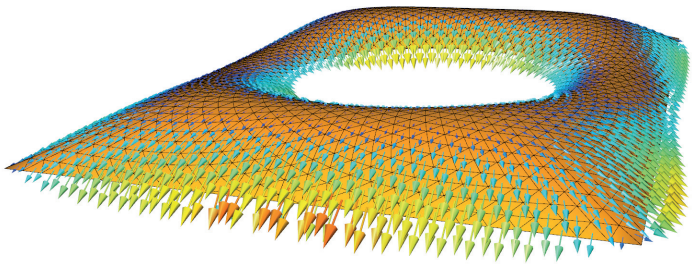

Figure 5: Example from Figure 3 continued: The Green-Lagrange tensor is uniformly positive-definite symmetric and can thus be brought to diagonal form with respect to its orthogonal eigenvector basis. Shear strains vanish in this particular basis so that the directions of minimal respectively maximal normal strain can be determined by sorting of eigenvalues. The latter serve as guidance field during initialization of the actual remeshing routine.

eigenvalues. These are called *principle strains* and equal the maximal respectively minimal absolute normal strain value encountered in a surface point.

A little care has to be taken concerning its orientation: If $\boldsymbol{e}_{\max}$ is an eigenvector of $\boldsymbol{\varepsilon}_{\mathrm{m}}$ corresponding to the eigenvalue greatest in modulus, then so is $-\boldsymbol{e}_{\max}$. A continuous sign distribution can be achieved, though, by aligning the guidance vector field $\boldsymbol{d} : \Gamma \mapsto T\Gamma$ with the projection of displacements $\boldsymbol{v}$ onto the tangent bundle $T\Gamma$ of Γ:

$$\boldsymbol{d} := \begin{cases} \boldsymbol{e}_{\max} & \text{if } \langle \boldsymbol{e}_{\max}, \mathbf{P}\boldsymbol{v} \rangle > 0, \\ -\boldsymbol{e}_{\max} & \text{else.} \end{cases}$$

Let us point out a significant strength of this specific choice of $\boldsymbol{d}$: Most quadrilateral meshing methods align isoparametric curves with one of the principle curvature directions. In umbilic areas, however, the latter are not well-defined and thus, some sort of smooth continuation is required. The mechanically motivated directions in turn are defined densely on the entire surface except for sets of zero measure, see Figure 5.

3 PQ Meshing Step

A typical output of the statics-sensitive preprocessing is illustrated in Figure 6(a), where integral curves of the guidance vector field are shown. Given such an output, we want to compute a PQ mesh fulfilling the following requirements:

- One family of mesh polygons shall be closely aligned with the guidance vector field,
- resulting quads shall be rather equally sized, and

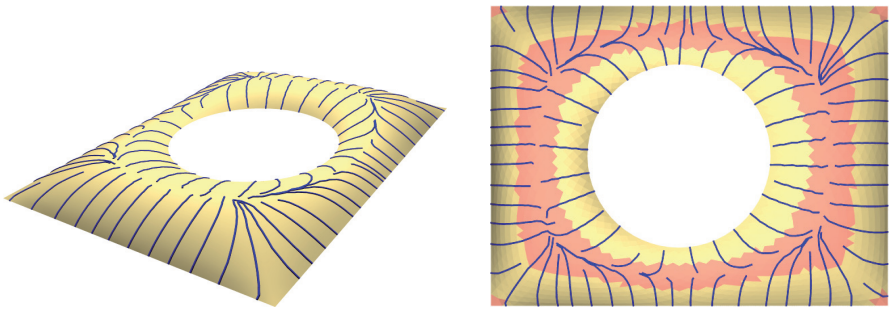

(a) (b)

Figure 6: Integral curves of the guidance vector field. The structure with fixed boundary has been subject to dead loading. It can clearly be seen that these integral curves are not suitable for quad remeshing: While they are nearly equally spaced along the boundaries, they become very close to each other and eventually converge to four points (the sources of the guidance vector field) in the interior. When laying out a quad mesh, we therefore do not care about alignment to the vector field in regions with relatively low maximum principal strain. Those regions are colored in brown in (b). The example is a model of the Great Court Roof by Foster and Partners, British Museum, London, compare Figure 1.

- resulting quads shall be planar within close tolerances.

Looking at the example in Figure 6(a), it becomes immediately clear that a resulting PQ mesh can not fulfill all of the above requirements at the same time: While the integral curves shown are nearly equally spaced along the boundaries, they become very close to each other and eventually converge to four points (the sources of the guidance vector field) in the interior. This behaviour depends on the specific load case and boundary conditions: In the example, the boundary is fixed, while gravitation imposes a uniform force distribution in the direction of the negative z-axis. Clearly, the deformation is maximal in the very interior of the surface. As strain arises from deformation by one differentiation, these regions are exactly where the zeros or, more precisely, the sources of the guidance vector field are located.

In order to find a reasonable solution fulfilling the requirements, we therefore relax the first one. We use the absolute maximum principal strain as a measure for the importance of alignment with the maximum loading directions and do not care about alignment in regions of the surface where it is below a predefined threshold. This is illustrated in Figure 6(a).

3.1 Level Set Method and Quad Meshing

In the following, we give an overview on how we create a quad-dominant mesh (a mesh consisting mainly of quads, except at the boundaries), which exhibits a family of edge polygons that is closely aligned with the guidance vector field, and whose quads are close to planar. We refer to [Zadravec et al. 2010] for the mathematical

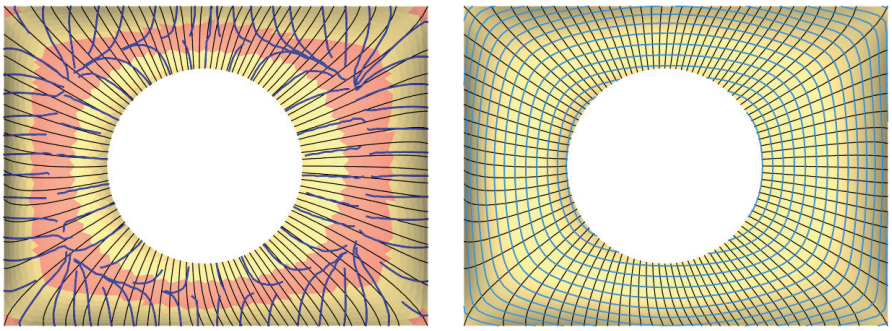

(a) (b)

Figure 7: (a) Integral curves of the guidance vector field as shown in Figure 6(a) overlayed with level set curves (black) resulting from an optimization for approximate alignment with the guidance vector field as well as equal spacing. It is easy to recognize that there is a tradeoff between these two optimization goals. In the brown region, alignment was not taken care of. (b) Transverse level set curves (blue) which have been optimized to become conjugate to the aligned level set curves as well as equally spaced. All curves are used for extraction of the quad-dominant mesh.

details. This remeshing task can be decomposed into the following steps:

1. **Edge polygons aligned with guidance vector field.** In this step we compute a family of curves on the reference surface that will later serve as input to the quad meshing. A straightforward way to compute curves that are perfectly aligned with the guidance vector field is shown in Figure 6(a): We could simply choose a set of starting points (e.g. equally spaced on the boundary as shown there) and trace integral curves of the guidance vector field from them. This would not allow us to control equal spacing of the curves. Therefore, we choose the following approach which was introduced in [Zadravec et al. 2010]: We compute a function defined on the reference surface by means of optimization such that its level sets (curves of constant function value) become equally spaced and closely aligned with the guidance vector field. We formulate these requirements as optimization goals in a least-squares sense, i.e. we do not require the goals to be fulfilled exactly but rather make use of the possibility to trade off between them. Moreover, we give up alignment with the guidance vector field in parts of the surface as described above. Figure 7(a) shows an example. We do not give details of the optimization here but refer the reader to [Zadravec et al. 2010].

2. **Edge polygons aligned with conjugate directions to guidance vector field.** There is a close relation between PQ meshes and conjugate curve networks on smooth surfaces (refer to [Liu et al. 2006] and [Zadravec et al. 2010]). Therefore, we aim at aligning the second family of edge polygons to curves

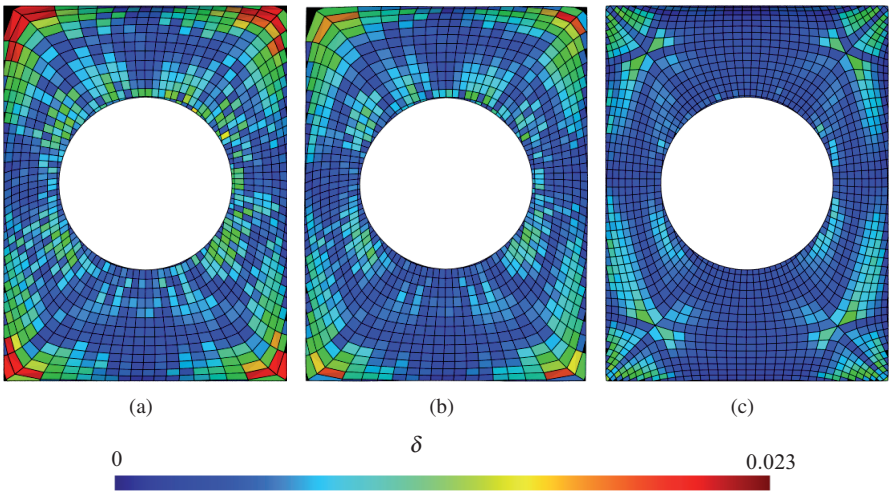

Figure 8: The Great Court Roof quad mesh (a) before and (b) after planarization. Planarity of quads is shown rainbow color coded: Blue corresponds to exact planarity, red corresponds to $\delta = 0.023$. For comparison, a conical PQ mesh is shown in (c). For a definition of conical mesh, see [Liu et al. 2006]. They offer the best possible choice for approximating a smooth surface by a PQ mesh, which is reflected by a low maximum $\delta = 0.015$. The connectivity is determined by the network of principal curvature lines and in general can not be used for our approach.

that are conjugate to the level set curves which we computed in the first step. We do this by computing a vector field that is conjugate to the guidance vector field. Then, we again employ the level set method as described above to compute a function whose level sets are aligned with this vector field. An example is shown in Figure 7(b).

3. **Quad meshing.** We use the level set curves computed using above steps to generate a quad mesh. We do this by intersecting the curves with each other and the boundary curves, and connecting neighboring intersection points. The resulting mesh will exhibit mostly quadrilateral faces, except at the boundaries as can be seen in Figures 9 and 13.

3.2 Planarization

The faces of meshes resulting from above procedure can be expected to be nearly planar because we use conjugate curve networks to generate them. This is illustrated by Figures 8(a) and 13(a). Nevertheless, we improve the planarity of the faces by subjecting the mesh to an optimization similar to [Liu et al. 2006], which slightly perturbs the positions of the vertices such that the faces become planar. Mathematical details can be found in [Zadravec et al. 2010]. The quality measure δ, used in

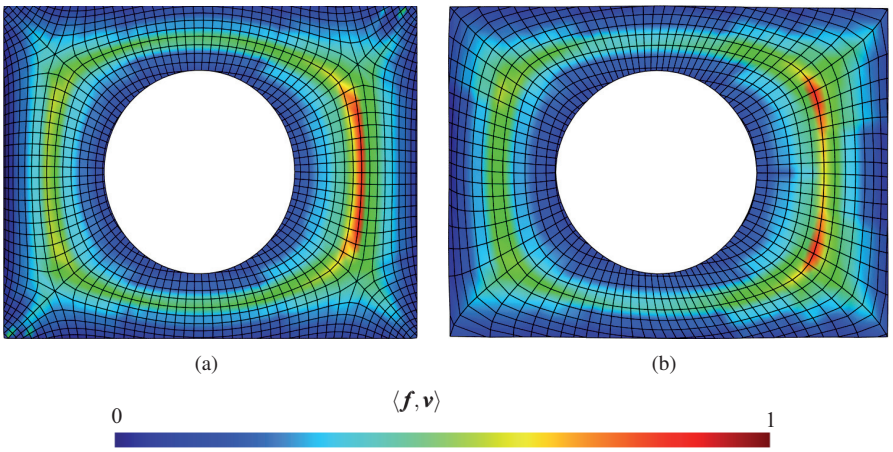

Figure 9: Compliance of the statically-optimized grid shell (b) is about 8 % lower than of the reference structure (a) .

the following for planarity of quad faces, corresponds to the distance between the diagonals of a quad divided by the mean length of its diagonals. For n-gons with $n > 4$, we consider planarity for all possible quads that are spanned by the vertices of the n-gon.

4 Results

Above all, the following discussion should assess the resulting quad meshes with respect to their mechanical properties. Particularly, we would like to oppose those which underwent the first phase of our algorithm to those which did not. A suitable performance measure is the local compliance, which is – as the scalar product of force and displacement fields $\langle \boldsymbol{f}, \boldsymbol{v} \rangle$ – inversely proportional to stiffness. The lower the compliance of a structure, the more favorable its statical behavior will be for a certain load case. Primarily, we must pass from the continuous viewpoint, i.e. from the microscopic triangle mesh serving as reference, to the discrete edge set of the resulting PQ mesh. As the former is assumed to be a membrane shell and thus non-resistant to bending, it translates to a *truss* in the discrete domain. Deformations of such derive from a similar kind of equilibrium condition as (1). Again, the system adopts the state of minimal potential $E_{\text{el}} - W$. In contrast to above, E_{el} now equals the (weighted) sum of elastic energies stored in every separate rod caused by extension/compression along its longitudinal axis. Let us remark that the discrete surface structures are sometimes called *grid shells*. This naming nicely illustrates their continuous origin.

4.1 Great Court Roof of British Museum

The first example geometry has already been introduced above. PQ approximations resembling the network of principal curvature lines are shown in Figures 8(b) (example taken from [Zadravec et al. 2010]) and 8(c). They differ mainly in the presence of four *singular vertices* or *branch points*, where more than four quads meet. This concept of *branched coverings* for improving the alignment with the guidance vector field is introduced in [Kälberer et al. 2007]. Nevertheless in this example, there is a tradeoff between the requirements as illustrated in Figure 7(a).

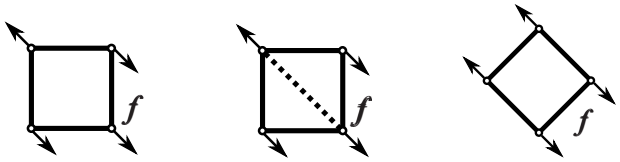

(a) High twisting moment (b) Stiffening by triangulation (c) Torsion-free alignment

Figure 10: Quadrilateral frame.

A compliance reduction becomes notable, mostly in the central regions where the deformation is highest (in modulus), as can be seen in Figure 9. In average, the reduction is about 8.0 %. The following supplies a heuristic explanation for why this should be so: Consider one grid cell in the discrete domain. Obviously, such a quadrangular frame is per se not very resistant to torsion, see Figure 10(a). In practice, supporting beams are inserted between two opposing corners to increase torsion stiffness, see Figure 10(b). Geometrically speaking, this would mean reverting to a triangular mesh. The sole remaining alternative is to orient the frame according to the known load case as indicated in Figure 10(c). This is roughly what our algorithm does, however, in a more abstract way and starting from a continuous perspective on the problem.

4.2 Neumuenster Abbey Court Roof Study

The given surface was originally designed by means of a combined structural and geometrical form finding process. Figure 11 shows that it was possible to fulfill the alignment with the guidance vector field very well while keeping the variation of spacing within tolerable bounds. We make the important observation that, like other e.g. curvature-based guidance fields, also the directions of principle strain may require the introduction of branch points into the quad mesh. In Figure 12, they appear where curves of maximal and minimal strain interchange roles. Figure 13 confirms that sufficient planarity was achieved.

We use a PQ mesh already shown in [Zadravec et al. 2010] for comparison of compliances. The impression one gets from Figures 14(a) and (b) is that the compliance *distribution* is less desirable in the mechanically motivated result. Still as predicted, its peak value of 0.8 is around 9 % lower than 0.88 in the reference model.

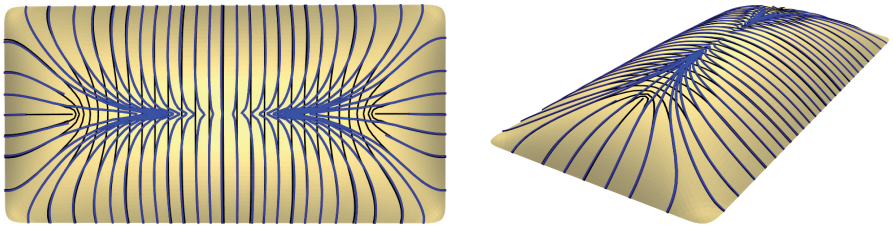

(a) (b)

Figure 11: Integral curves of the guidance vector field (blue) and resulting level set curves (black) for the Neumuenster Abbey Court Roof example. Loading and boundary conditions are the same as in Section 4.1. In contrast to the example shown in Figure 6(a), the guidance vector field exhibits only one source. The level set curves are aligned very well to the guidance vector field except in the central region, where maximum principal strains are lower and integral curves would obviously not be usable for quad meshing.

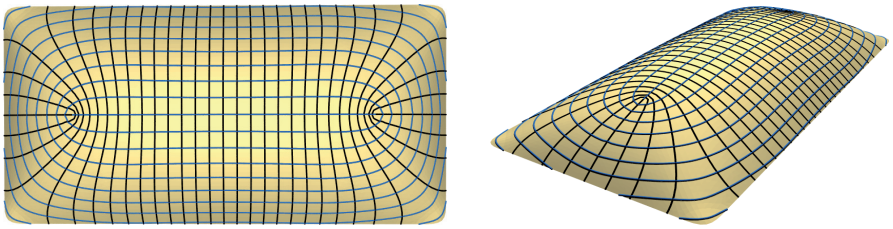

(a) (b)

Figure 12: Level set curves aligned with the guidance vector field (black) and transverse level set curves (blue) optimized to become equally spaced and conjugate to the black curves. Quad meshing has been done based on these curves, see Figure 13.

5 Conclusion

We put forward a method for the layout of quadrilateral meshes which are almost planar and, beyond that, optimized with respect to their statics. Therein, physical and geometric computations are decoupled as stringently as possible to reduce conflicts among the multiplicity of involved objectives and constraints: Using a shell model of the reference surface, a first step yields the direction of maximal principle strain in every point. Then, a not necessarily planar quad-dominant mesh is generated with one family of isoparameter lines integrating the maximal principal strain vector field wherever possible. The algorithm concludes with an established planarization procedure, preserving the initial mechanical properties in large part. A compliance analysis suggests superiority of the statically optimized layouts over those arising from purely geometrical reasoning.

One future extension of the presented framework is obvious: While the basic

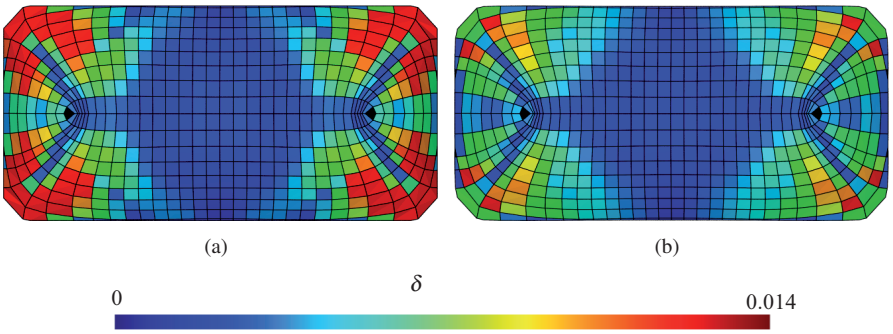

Figure 13: The Neumuenster Abbey Court Roof quad mesh (a) before and (b) after planarization. Planarity of quads is shown rainbow color coded: Blue corresponds to exact planarity, red corresponds to $\delta = 0.014$.

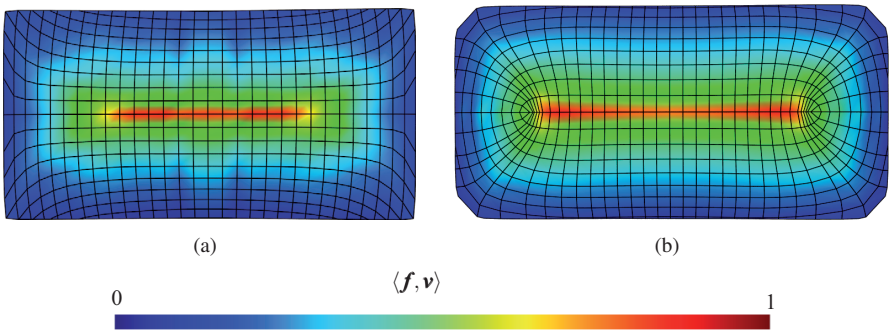

Figure 14: Remeshing result (a) with and (b) without mechanical preprocessing.

principle is well demonstrated under a membrane assumption, only a FEM implementation suitable for fourth-order PDEs would accommodate sufficiently general deformation modes like strong bending. Moreover, we plan to incorporate a shape optimization step, in which now the geometry itself is treated as variable – of course within tolerances describing the architect's willingness to accept modifications of the original design.

Acknowledgments

The research leading to these results has received funding from the European Communitys Seventh Framework Programme under grant agreement No. 230520 (ARC), and grant No. 813391 of the Austrian research promotion agency (FFG). The authors would like to thank Waagner-Biro for providing the digital model of

the British Museum courtyard roof, and RFR for providing the Neumuenster Abbey court roof mesh.

References

BOTSCH, M., AND SORKINE, O. 2008. On linear variational surface deformation methods. *IEEE Transactions on Visualization and Computer Graphics 14*, 1, 213–230.

CIRAK, F., ORTIZ, M., AND SCHRÖDER, P. 2000. Subdivision surfaces: a new paradigm for thin-shell finite-element analysis. *International Journal for Numerical Methods in Engineering 47*, 12, 2039–2072.

DAVIS, T., AND HAGER, W. 2009. Dynamic supernodes in sparse Cholesky update/downdate and triangular solves. *ACM Transactions on Mathematical Software 35*, 4, 1–27.

GLYMPH, J., SHELDEN, D., CECCATO, C., MUSSEL, J., AND SCHOBER, H. 2002. A parametric strategy for freeform glass structures using quadrilateral planar facets. In *Acadia 2002*, ACM, 303–321.

GRINSPUN, E., HIRANI, A., DESBRUN, M., AND SCHRÖDER., P. 2003. Discrete shells. In *Proc. of ACM SIGGRAPH/Eurographics Symposium on Computer Animation*.

HUGHES, T., COTTRELL, J., AND BAZILEVS, Y. 2005. Isogeometric analysis: CAD, finite elements, nurbs, exact geometry and mesh refinement. *Computational Methods in Applied Mechanical Engineering 194*, 4135–4195.

HUGHES, T. 2000. *The Finite Element Method*. Dover Publications.

KÄLBERER, F., NIESER, M., AND POLTHIER, K. 2007. QuadCover – surface parameterization using branched coverings. *Computer Graphics Forum 26*, 3, 375–384. Proc. Eurographics.

KOITER, W. 1970. On the foundations of the linear theory of thin elastic shells. In *Proc. Kon. Ned. Ak. Wet.*, 169–195.

LIU, Y., POTTMANN, H., WALLNER, J., YANG, Y. L., AND WANG, W. 2006. Geometric modeling with conical meshes and developable surfaces. *ACM Trans. Graphics 25*, 3, 681–689.

LOVE, A. 1888. On the small vibrations and deformations of thin elastic shells. *Philosophical Transactions of the Royal Society*, 179–491.

NAGHDI, P. 1972. *Handbuch der Physik*. Springer, ch. The theory of shells and plates, 425–640.

REISSNER, E. 1947. On bending of elastic plates. *Quart. Appl. Math. 55*, 45–68.

WINSLOW, P., PELLEGRINO, S., AND SHARMA, S. 2008. Multi-objective optimization of free-form grid structures. In *Proceedings of the International Conference on Engineering Optimization*.

ZADRAVEC, M., SCHIFTNER, A., AND WALLNER, J. 2010. Designing quad-dominant meshes with planar faces. vol. 29. Proc. Symp. Geometry Processing.

Authors' addresses:

Alexander Schiftner (aschiftner@geometrie.tuwien.ac.at):
Technische Universität Wien
Wiedner Hauptstr. 8–10/104
1040 Wien, Austria.

Jonathan Balzer (jonathan.balzer@kaust.edu.sa):
King Abdullah University of Science and Technology
Center for Geometric Modeling and Scientific Visualization
Thuwal 23955-6900
Kingdom of Saudi Arabia.