

QUAD LAYOUTS WITH HIGH VALENCE SINGULARITIES FOR FLEXIBLE QUAD MESHING

Jovana Jezdimirović¹ Alexandre Chemin¹ Maxence Reberol¹ François Henrotte¹
Jean-François Remacle¹

¹*Université catholique de Louvain, Louvain la Neuve, Belgium. jovana.jezdimirovic@uclouvain.be*

ABSTRACT

A novel algorithm that produces a quad layout based on an imposed set of singularities is proposed. In this paper, we either use singularities that appear naturally, e.g., by minimizing Ginzburg-Landau energy, or use as an input user-defined singularity pattern, possibly with high valence singularities that do not appear naturally in cross-field computations. The first contribution of the paper is the development of a formulation that allows computing a cross-field from a given set of singularities through the resolution of two linear PDEs. Specific mesh refinement is applied at the vicinity of singularities to accommodate the large gradients of cross directions that appear in the vicinity of singularities of high valence. The paper's second contribution is a correction scheme that repairs limit cycles and/or non-quadrilateral patches. Finally, a high-quality (elements edges being as close to orthogonal as possible) block-structured quad mesh is generated from the quad layout and per-partition parametrization.

Keywords: quad layout, quad meshing, high valence singularities

1. INTRODUCTION

Quad meshing is a discipline that has been developed for different purposes by two distinct communities. Computer graphics and engineering analysis communities have indeed produced extensive research in the subject for decades, leading to both unstructured and structured approaches of quad meshing. A recent survey [1] comprehensively discusses the issue of quad meshing from manual to fully automatic generation, along with the respective advantages and drawbacks of the different methods.

In a block-structured quad mesh, most of the vertices are regular. A vertex v is regular if it has the optimal number n_v of adjacent quads: $n_v = 2$ if the vertex is on the boundary $\partial\Omega$ of the domain Ω to be meshed, and $n_v = 4$ if it is an internal vertex. Our approach for block-structured quad meshing relies on the computation of an auxiliary object, called a cross-field, as a guide to orient the quadrangular elements. Cross-fields have isolated singularities that can be shown to

correspond to the location of irregular vertices of the quad mesh [2]. Using this useful structure of cross-fields has become popular to generate block-structured quad meshes [3, 4, 5, 6, 7]. More specifically, generating block-structured quad meshes requires the construction of a *quad layout*. A quad layout is the partitioning of an object's surface into simple networks of conforming quadrilateral patches [8]. Fig. 1 shows examples of quad layouts for a simple 2D domain.

Two main approaches exist to extract a quad layout from the singular structure of a cross-field. The first approach consists in computing a seamless global parametrization of the domain where integer iso-values of the parameter fields form the sides, or the separatrices, of the quad layout [9, 10, 11]. This approach is implicit in the sense that it does not explicitly compute separatrices from the cross-field. The second approach is explicit. It is the one used in this paper. It consists in directly tracing the separatrices of the quad layout, starting at the singularities of the cross-field. A singularity of valence 5 corresponds to a vertex v with




Figure 1: Three quad layouts of a simple domain defined by a square minus a disk. Colored points are the irregular corners of the layouts.

$n_v = 5$ adjacent quads, so that 5 separatrices can be started at this vertex v . Separatrices then follow integral lines of the cross-field, eventually ending on $\partial\Omega$ or on another singularity.

The main drawback of the explicit approach is related to its lack of geometrical robustness. Assuming that singularities are located at vertices of our mesh, the orientation of the cross-field is thus varying abruptly at the vicinity of a singular vertex v . That bad representation of the cross-field can have dramatic consequences. In many cases, the right number of separatrices cannot even be traced starting at v . Bad cross representation around singularities is also the cause of tangential crossings, i.e., separatrices that should meet at a singular vertex but that miss the rendez-vous and continue further, eventually forming an infinite loop.

Mesh refinement around singularities would then seem to be a reasonable idea. Disappointingly, it is not. Cross-field smoothers have the bad habit to push singularities away from refined zones, because the energies used in cross-field models are unbounded (in the continuous case) at singularities, so that locating singularities in coarse region allows finding better minima (in the discrete case). Standard mesh adaptation thus fails here, because the solution does not converge in terms of energy. Fig. 2 shows an example of that typical behavior. Fig. 2 (left) shows a cross-field computed on a uniform mesh. Mesh refinement is then applied close to those singularities, and the cross-field solver is then applied once again with the refined mesh. On Fig. 2 (right), one can observe that the singularities have moved with respect to their position after the first cross-field computation. The tendency to escape refined regions is so strong that the solution is actually breaking the symmetry of the problem.

In order to avoid this issue, our cross-field computation is implemented as a two-step process. A first cross-field is computed in a standard fashion by means of minimizing a nonlinear energy (the Ginzburg-Landau

energy functional in our case). The *singularity pattern* of the cross-field is extracted from this first computation and then injected as a constraint for a second cross-field computation on an adapted mesh. The formulation for computing this second cross-field is different. Here, the singularity pattern is taken as an input and a cross-field is computed by means of solving two linear systems. Singularity positions *are thus prescribed* so mesh adaptation can be performed. The local structure of a cross-field at a singularity is essentially radial: the mesh that is used in this second step thus involves specific *bicycle spokes patterns* at singularities (see section 4). With that *ad hoc* refinement, the singular structure of the cross-field is perfectly captured. Another advantage of our approach is the possibility of moving singularities, adding or removing some, while still respecting topological constraints. In this work, we use that advantage for correcting or adapting the quad layout when necessary, e.g., to fix non-quadrilateral partitions or to avoid limit cycles.

The full pipeline of our block-structured quad mesher consists in four steps and is described in Figure 3.

Step 1: compute a first cross-field using standard non linear minimization procedure that allows to find a singularity pattern i.e. singularity positions and indices.

Step 2: compute a second cross-field with the prescribed singularity pattern of Step 1 on a refined mesh.

Step 3: compute a neat quad layout decomposition without infinite loops or t-junctions. Quad layout partitioning is performed on the accurate cross-field of Step 2 by applying the separatrix tracing scheme presented in [7].

Step 4: generate a full quadrilateral mesh.

Note that Step 1 can be skipped if a singularity pattern is provided by the user. Figure 4 shows this reduced 3-step process with a very unique singularity of very high valence 8!




Figure 2: Displacement of singular points due to a cross-field mesh dependency. Cross-field computation follows the work from [2].

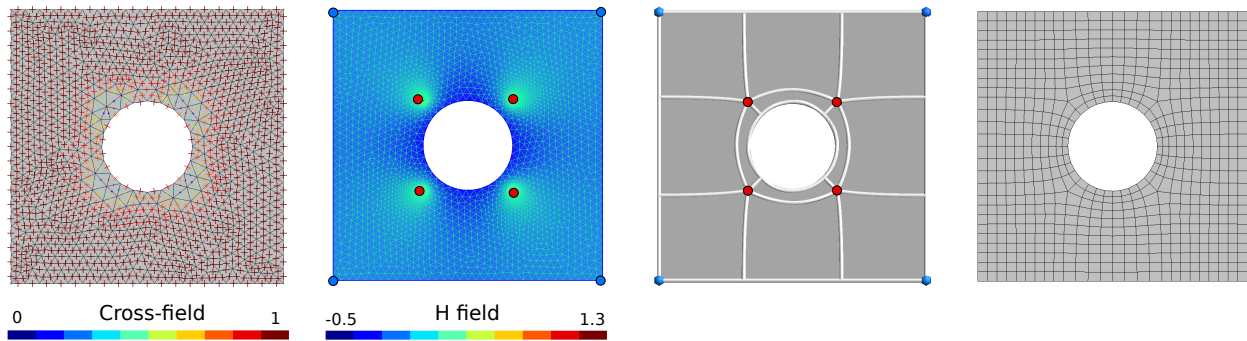


Figure 3: Illustration of 4-steps pipeline




Figure 4: Illustration of the reduced 3-steps pipeline

The two-step cross-field computation delivers thus an accurate mesh independent cross-field, Fig. 3, from which one can then continue with a neat quad layout decomposition, and eventually a full quadrilateral mesh, Fig. 3 and 4. Quad layout partitioning is performed on the final (accurate) cross-field by applying the separatrix tracing scheme presented in [7]. The

final step is the generation of the quadrilateral mesh by following the parametrization imposed by the cross-field (section 7).

The steps of the proposed approach are detailed in Algorithm 1. The whole pipeline has been implemented as a fully functional module in Gmsh [12], the open

source finite element mesh generator. Although the pipeline can be fully automated, a certain level of interactivity is possible. The user is indeed given the opportunity to modify the singularity pattern and to adapt the location and index of singularities, as long as it corresponds to the topology of the domain. The modified singularity valence can be as high as six or eight, Fig. 4.

Algorithm 1 Generate all quad mesh

INPUT: FILE in `.geo` format

OUTPUT: FILE in `.msh` format

- 1: Non-linear cross-field computation on Ω (Sect. 3)
 - 2: Extract singularity pattern (Sect. 4)
 - 3: Compute cross-field with imposed singularities (Sect. 5)
 - 4: Obtain and correct quad layout (Sect. 6)
 - 5: Generate quad mesh on Ω (Sect. 7)
-

2. RELATED WORK

Due to the fact that some of the techniques used in our approach are customary in the meshing and computational geometry communities, we will focus in this article on the most notable trends and contrasts/resemblances with existing works. For the sake of comprehensiveness, bibliographic references for our algorithm are indicated in corresponding sections.

2.1 Cross-field generation for meshing purposes

For quad meshing purposes, authors are generally relying on either explicit quadrangulations or parametrization/remeshing techniques [13]. The group of parametrization approaches is enriched with numerous methods [1], where rough dissemination could be made with two generalized flows: cross-field guidance algorithms ([4, 5, 6, 7]) and cut-graph techniques ([9, 10]). Our approach follows the first flow with the difference of using two independent successive cross-fields, which is unusual in related work on cross-field generation ([14, 15]). Similarly to [7, 11, 16], we aim at obtaining a locally integrable cross-fields, such that for each point X inside a patch holds:

$$\nabla \times \mathbf{u}_1 = 0, \quad \nabla \times \mathbf{u}_2 = 0, \quad (1)$$

where \mathbf{u}_1 and \mathbf{u}_2 represent two orthogonal branches of the cross defined in point \mathbf{X} .

Stressed out by [17], a major compromise to be done in the computation of cross-fields is between efficient design of vector field and control over singularities. Our algorithm runs fully automatic which frees the user of painful task to explicitly specify all singularities

on complex domains [18]. At the same time, due to the independent steps of the algorithm, we are also able to fulfill the aim of generating a fast smooth cross-field that respects a user-defined singularity patterns as in [13, 14].

2.2 Singularities with high indices

High-order singularities' patterns play an important role in numerous scientific applications ([19, 20, 21, 22, 23]). Initially, high-order singularities' definition, extraction and visualization were relying on mathematical background of Clifford algebra, typically motivated by applications for fluid mechanics or electrostatics ([19, 22, 23]). Some more recent papers ([21, 23]) are focused on their role in vector field generation. In our case, the application of high-order singular points is exploited for its benefits for quad meshing purposes, specifically obtaining a quad layout and inducing important element size gradients in the refined block-structured quad mesh.

When it comes to imposed singularity configuration, it is essential to note that not all of the topologically correct singularity patterns will be valid for quad meshing. The work [24] presents the sufficient and necessary condition for valid singularity configuration, where both the singularities' valences and holonomy are taken into account. Some of the most interesting findings on the non-existence of conformal quadrangulations can be found in [24, 25, 26, 27].

2.3 Quad layout partition guided by cross-field

Representation of geometrically or topologically complex domains requires their decomposition into blocks, preferably quadrilateral. A wide range of existing methods is described in the survey [1]. Works analogous to ours, i.e. cross-field based approaches, are decomposing the domain either by generating separatrix graph obtained from numerical integration of streamlines ([4, 5, 6, 7, 28]) or by constructing edge maps ([29, 30, 31]) and motorcycle graph ([32]) using the isolines of an underlying parameterization. As already reported by authors working on the first group of approaches, generated quad layout commonly requires repairing of invalid configurations caused by the appearance of limit cycles and thin blocks. Some of the efficient methods for solving these issues are [29, 28, 33, 34] as well as the quantization method [9]. Generating and correcting the quad layout, in our case, shares some common points with [7, 9, 28], exact differences with respect to these works will be outlined in the corresponding sections.

3. FIRST CROSS-FIELD COMPUTATION

A 2D cross \mathbf{c} is a set of 4 orthogonal vectors of the same norm l , i.e., $\mathbf{c} = \{\mathbf{u}_k, |\mathbf{u}_k| = 1, k \in [1, 4]\}$. On a planar 2D domain Ω , a Cartesian basis $\{\mathbf{x}, \mathbf{y}\}$ can be defined, and the angular orientation of the cross \mathbf{c} can be simply represented by an angle θ in that basis, since one has $\mathbf{u}_k = l \cos(\theta + k\frac{\pi}{2})\mathbf{x} + l \sin(\theta + k\frac{\pi}{2})\mathbf{y}$, as depicted in Fig. 5. This angular representation θ of the cross is however determined only up to an additive factor $k'\frac{\pi}{2}$ with $k' \in \mathbb{Z}$.




Figure 5: 2D cross definition as a set of four orthogonal vectors

A 2D cross-field, now, is a map $\mathcal{C}_\Omega : \mathbf{X} \in \Omega \rightarrow \mathbf{c}(\mathbf{X})$, and the standard approach to compute a smooth boundary-aligned cross-field is to minimize the Dirichlet energy

$$\min_{\mathcal{C}_\Omega} \int_{\Omega} \|\nabla \mathcal{C}_\Omega\|^2 \quad (2)$$

subject to the boundary condition $\mathbf{c}(\mathbf{X}) = \mathbf{g}(\mathbf{X})$ on $\partial\Omega$, where \mathbf{g} is a given function.

The indetermination on θ mentioned above prevents from using it directly as unknown for the finite element formulation of the problem. In order to have a unique (symmetry invariant) and continuous cross representation, the alternative vector representation $\mathbf{v}(\mathbf{X}) = (\cos(4\theta), \sin(4\theta))$ is preferred [35], which leaves us with a vector field to solve for.

This vector field is continuous, but also subject to the nonlinear constraint $|\mathbf{v}(\mathbf{X})| = 1$. An elegant way to deal with this constraint is to minimize a Ginzburg-Landau type functional, which penalizes local vector values that deviate from the unit norm [2]. Direct solving of the non-linear Ginzburg-Landau problem is however rather expensive. Another approach is the MBO algorithm proposed in [6]. It is a faster and simpler alternative that provably converges to the same results on uniform meshes. The main idea of the MBO algorithm is to alternate a heat diffusion problem, that performs the Laplacian smoothing of the vector field, with a projection back onto the space of unit norm cross-fields, until convergence is achieved. Each step

of MBO algorithm is thus given by

$$\frac{\partial \mathbf{v}}{\partial t} = \alpha \nabla^2 \mathbf{v} \quad \text{then} \quad \mathbf{v} \leftarrow \frac{\mathbf{v}}{\|\mathbf{v}\|}. \quad (3)$$

As in [36], we start with large diffusion (e.g. $\alpha = (0.1 * \text{bbox diagonal})^2$) and we progressively decrease until we reach the small diffusivity coefficient associated to the mesh size (e.g. $\alpha = (\text{edge min})^2$). This strategy considerably accelerates convergence as the low frequencies are solved before the larger ones, as in a multi-grid approach. To further accelerate the computation, we use a small number of diffusivity coefficients (e.g., five to ten) and we apply the MBO step until convergence for each diffusivity coefficient. As the diffusion linear system is not changing for a given diffusivity coefficient, we can factorize the linear system one time and reuse it until convergence. In the end, the computational cost is approximately one linear system solve per diffusivity coefficient, i.e., five to ten in total, which is very reasonable to accurately solve the cross-field non-linear problem (Eq. 2).

It is important to note that the aim of this first cross-field computation is to locate singularities. Therefore, any cross-field computation method could be used here instead, as for example the one described in [37], which could prove to be faster.

4. SINGULARITY PATTERN

Singularities play an important role for meshing purposes and their properties are well documented in the literature [4, 5, 6, 7, 28]. Obviously, cross-fields have a close parenthood with vector fields and, like any direction field (a normalized vector field, a cross-field, . . . , more details in [14]), they are subject to the topological constraints of the Poincaré-Hopf theorem. Let us consider, for instance, a smooth unit norm vector field \mathbf{u} on a surface Ω . The local value of this vector field can be defined as the rotation by a certain angle θ with respect to a local angular reference around the local normal vector to the surface. This natural geometric representation of the vector field has however a number of issues. It is only defined up to an additive angle $2k\pi, k \in \mathbb{Z}$. It is also discontinuous wherever the value of θ jumps from 0 to 2π or reversely. Finally, as customary in differential geometry, θ values at distant points are incommensurable unless a parallel transport rule is added, because they refer to different local normal vectors and different local reference positions.

Still, in an infinitesimal neighborhood V_X of a point $X \in \Omega$, the parallel transport is superfluous because both the normal vector and the reference angular reference can be considered constant. The smooth vector field \mathbf{u} then induces a map $C^1 \rightarrow C^1$, where C^1 is the unit circle, for any (infinitesimal) closed curve

$C \subset V_X$. This map follows the turns done by \mathbf{u} as the curve C is travelled once. As \mathbf{u} is smooth, the number of turns is necessarily an integer number, called the *index* of \mathbf{u} at X , $\text{index}_X(\mathbf{u}) \in \mathbb{Z}$. The point X is regular if the index is zero, and it is a singularity of index $\text{index}_X(\mathbf{u})$ otherwise.

The principle is similar for a cross-field \mathcal{C}_Ω except that the map is $C^1 \rightarrow C^1/O$ where O stands for the symmetry group of the cross, and that the index is now an integer multiple of $\frac{1}{4}$. If one notes $S_j, j = 1, \dots, N$, the singularities of the cross-field, the Poincaré-Hopf theorem states that

$$\sum_{j=1}^N \text{index}_{\mathcal{C}_\Omega}(S_j) = \chi(\Omega), \quad (4)$$

where $\chi(\Omega)$ is the Euler characteristic of the (non necessarily planar) surface Ω . The indices of the singularities will be noted

$$k_j \equiv \text{index}_{\mathcal{C}_\Omega}(S_j) = \frac{t}{4}, \quad t \in \mathbb{Z} \quad (5)$$

and we shall also use the related concept of valence defined by

$$k_j = \frac{4 - \text{valence}(S_j)}{4}, \quad \text{valence}(S_j) = 4 - 4k_j.$$

In practice, we use a simple way to find the singularity location by computing the winding number in a similar fashion as in [28], originally from [38]. The corresponding valence is then determined by detecting the number of separatrices concurring at the singularity [7]. When the cross-field is computed on a domain with a coarse mesh around singularities (of valence five or higher), numerical inaccuracies can be expected in the determination of the cross-field orientations, which will finally lead to an invalid quad layout generation. To prevent this, the mesh is refined in the form of bicycle spokes in the vicinity of identified singularities, so that the large gradients of cross directions are well accommodated, as shown in Fig 6.

An adaptation of the singularity pattern can be performed at this stage of the pipeline. With the two-step approach, the user is indeed offered the opportunity before the second cross-field computation to modify the location of singularities, or even to place additional singularities, provided the final singularity pattern still fulfills the Poincaré-Hopf topological constraint (4). For instance, additional pairs of singularities of valence 3 and 5 do not modify the result of the summation in (4). Such 3-5 pairs can therefore be added freely to the singularity pattern in order, for instance, to allow a less distorted quad layout or to impose a specific size map, as shown in Fig. 7.

In practice, imposing a singularity pattern simply consists in encoding valences data for vertices of the initial




Figure 6: Mesh refinement around singularity of valence 8




Figure 7: Two quad layouts of the same region. The addition of a 3-5 pair to the singularity pattern, in the second example, allows imposing a specific size map on the region.

mesh and ensuring a sufficient mesh refinement in their respective vicinities.

It will also be shown at the end of the following section that not all valid singularity patterns will be suitable for obtaining a high quality quad layout, depending on the geometrical characteristics of the domain. Mispositioning of singularities or imposing inadequate valences may lead in practice to non-quadrilateral patches.

5. CROSS-FIELD COMPUTATION WITH IMPOSED SINGULARITY PATTERN

The nonlinear problem described in Sect. 3 is able to reveal a singularity pattern for the domain Ω under analysis. The computed cross-field \mathbf{v} is however rather inaccurate. It is indeed continuous and singularities are areas in this continuous field where the norm deviation from unity fails to be ruled out by the penalty term. This inaccuracy does not heavily affect the singularity pattern, because singularity indices are integer quantities and their location need

only to be known approximately. However, it jeopardizes the accurate subsequent tracing of separatrices and thus justifies the construction of a second more accurate cross-field with a truly nonlinear field representation, and whose computation is based on the singularity pattern that has been extracted from the first computation (Sect. 4).

Let the singularity pattern be defined as the set $\mathcal{S} = \{S_j, j \in [1, N]\}$ of N singularities located at \mathbf{X}_j and of respective index k_j . A cross-field \mathcal{C}_Ω matching this singularity pattern must fulfill three constraints:

- if \mathbf{X} is on $\partial\Omega$, at least one branch of $\mathcal{C}_\Omega(\mathbf{X})$ is tangent to $\partial\Omega$,
- if \mathbf{X} is a singularity of \mathcal{C}_Ω then $\mathbf{X} \in \mathcal{S}$,
- if $\mathbf{X} = S_j$ then $\text{index}(\mathcal{C}_\Omega(\mathbf{X})) = k_j$.

As in this paper only the *planar case* is considered, $\Omega \subset \mathbb{C}$ can be regarded as a subset of the complex plane. We shall note \mathbf{n} the unit normal vector to the complex plane, \mathbf{g} the interior unit normal vector to the boundary $\partial\Omega$, and $\mathbf{T} = \mathbf{g} \times \mathbf{n}$ the unit vector tangent to $\partial\Omega$.

The method presented in this section relies on the computation of a dual scalar field H defined in section 5.1, continuous on the whole domain Ω . This approach is similar as the one detailed in [39].

5.1 Computation of the H field

We are looking for a holomorphic function

$$\mathcal{F} : \begin{array}{l} \mathcal{P} \subset \mathbb{C} \rightarrow \mathbb{C} \\ \mathbf{z} = U + iV \mapsto \mathbf{P} = x + iy \end{array} \quad (6)$$

where \mathcal{P} is a parametric space. As \mathcal{F} and \mathcal{P} are unknown, solving directly this problem is proved to be a difficult task. Instead, the focus is on finding the 2×2 jacobian matrix of \mathcal{F}

$$J_{\mathcal{F}}(\mathbf{z}) = (\partial_U \mathcal{F}(\mathbf{z}), \partial_V \mathcal{F}(\mathbf{z})) \equiv (\tilde{\mathbf{u}}, \tilde{\mathbf{v}}), \quad (7)$$

where $\tilde{\mathbf{u}}, \tilde{\mathbf{v}}$ are the column vectors of J . The function \mathcal{F} being holomorphic, the columns of J have the same norm $\|\tilde{\mathbf{u}}\| = \|\tilde{\mathbf{v}}\| \equiv e^H$ and are orthogonal to each other, $\tilde{\mathbf{u}} \cdot \tilde{\mathbf{v}} = 0$, so that one can write

$$J_{\mathcal{F}}(\mathbf{z}) = e^H (\mathbf{u}, \mathbf{v}) \quad , \quad \mathbf{n} = \mathbf{u} \wedge \mathbf{v}$$

$\tilde{\mathbf{u}}$ and $\tilde{\mathbf{v}}$ are 2 branches of the cross-field we are looking for, and $\tilde{\mathbf{u}}$ can be represented by a complex function

$$f : \Omega \rightarrow \mathbb{C} \quad , \quad f = e^{H+i\theta}. \quad (8)$$

The computation of the cross-field is rephrased into the computation of two real functions, $H : \Omega \rightarrow \mathbb{R}$ and

$\theta : \Omega \rightarrow \mathbb{R}$, from which the first branch $\tilde{\mathbf{u}} = f = e^{H+i\theta}$ can be deduced, and the second branch of the cross-field is simply $\tilde{\mathbf{v}} = i\tilde{\mathbf{u}} = if$.

By construction, $\tilde{\mathbf{u}}$ being a partial derivative of a holomorphic function, f is holomorphic. As (8) shows, H and θ are the real and imaginary parts of a complex logarithmic function. Indeed,

$$\begin{aligned} \log(f) &= \ln(|f|) + i \arg(f) \\ &= H + i(\theta + 2k\pi), k \in \mathbb{Z}. \end{aligned} \quad (9)$$

The complex logarithm being a holomorphic function, both H and θ are harmonic real functions. The H field, which is the real part of the logarithm, is continuous on $\Omega \setminus \mathcal{S}$, i.e., on the domain Ω from which the pointwise singularities \mathcal{S}_j have been excluded. The θ field, on the other hand, is multi-valued and a *branch cut* (see Sect. 5.2.1) has to be additionally defined on $\Omega \setminus \mathcal{S}$ to have it single-valued, and make its computation possible. Finally H and θ obey the Cauchy-Riemann equations

$$\nabla\theta = \mathbf{n} \times \nabla H, \quad (10)$$

which shall allow to obtain θ , once H has been computed.

The H field is solved first. As shown in [40], it is the solution of the partial differential problem on Ω

$$\nabla^2 H = 2\pi \sum_{j=1}^N \frac{k_j}{4} \delta(\mathbf{X}_j). \quad (11)$$

The boundary condition $\mathbf{T}(s) = (\cos \theta(s), \sin \theta(s))$ ensures that the cross, represented by $\theta(s)$, is tangent to the boundary $\partial\Omega$, represented by the tangent vector $\mathbf{T}(s)$, with s a curvilinear coordinate on $\partial\Omega$. This entails by derivation

$$\partial_s \mathbf{T} = (-\sin \theta, \cos \theta) \partial_s \theta = \mathbf{g} \partial_s \theta.$$

On the other hand, one knows from the differential geometry of surfaces that $\partial_s \mathbf{T} \equiv \kappa_g \mathbf{g}$ with κ_g the geodesic curvature of the curve $\partial\Omega$. One has thus by identification

$$\kappa_g = \partial_s \theta = \mathbf{T} \cdot \nabla \theta = (\mathbf{n} \times \nabla H) \cdot \mathbf{T} = -\mathbf{g} \cdot \nabla H,$$

where Eq. (10) has been used.

Summing up, the computation of H consists in solving the linear boundary value problem

$$\begin{cases} \nabla^2 H &= 2\pi \sum_{j=1}^N \frac{k_j}{4} \delta(\mathbf{X}_j) & \text{in } \Omega \\ -\mathbf{g} \cdot \nabla H &= \kappa_g & \text{on } \partial\Omega. \end{cases} \quad (12)$$

As the boundary condition is of Neumann type on the entire boundary, a solution H exists if the condition

$$\int_{\partial\Omega} \kappa_g \, d\partial\Omega = \int_{\Omega} 2\pi \sum_{j=1}^N \frac{k_j}{4} \delta(\mathbf{X}_j) \, d\Omega \quad (13)$$

is verified, which is obviously the case, as the Gauss-Bonnet theorem states the first identity in

$$\int_{\partial\Omega} \kappa_g \, d\partial\Omega = 2\pi\chi(\Omega) = \int_{\Omega} 2\pi \sum_{j=1}^N \frac{k_j}{4} \delta(\mathbf{X}_j) \, d\Omega \quad (14)$$

whereas the second identity is the Poincaré-Hopf theorem.

The solution H of the boundary value problem Eq. (12) is however known up to an arbitrary additive constant, which is not harmful as only ∇H is needed to determine θ . It is a linear problem solved using the finite element method on a triangulation of Ω with order 1 Lagrange elements. Once H is known, one proceeds with the determination of θ , as explained in the next section, to complete the cross-field $(f, \iota f)$.




Figure 8: H field obtained for 4 singularities of valence 3 on corners, and 4 singularities of valence 5 inside the domain.

5.2 Computation of the θ field

As the curl of the right-hand side of Eq. (10)

$$\nabla \times (\mathbf{n} \times \nabla H) = (\nabla^2 H) \mathbf{n} = \mathbf{0}$$

vanishes in $\Omega \setminus \mathcal{S}$, one knows that the left-hand side is indeed a gradient, and that there exists a scalar field θ verifying the differential equation (10). This function is however multi-valued and a branch cut has to be determined to make it single-valued, and make its computation possible.

5.2.1 Generate the branch cut

A branch cut is a set \mathcal{L} of curves of a domain Ω that do not form any closed loop and that cut the domain in such a way that it is impossible to find any closed loop in $\Omega \setminus \mathcal{L}$ that encloses one or several singularities, or an internal boundary. As we already have a triangulation of Ω , the branch cut \mathcal{L} is in practice simply a set of edges of the triangulation. The method presented here is based on [13].

The nodes and the edges of a triangulation of a domain Ω can be regarded as a graph \mathcal{G} . A spanning tree \mathcal{T} is a sub-graph of \mathcal{G} that links all vertices of \mathcal{G} with the minimum number of edges. A spanning tree can be generated by a Breadth-first-search algorithm. One starts from an arbitrary vertex v , and adds to the spanning tree \mathcal{T} all edges (v, v_i) that are adjacent to v and that connect vertices v_i that have not yet been visited by the algorithm. The procedure is continued recursively with all newly considered vertices v_i until running out of unvisited vertices. There exist many equivalent spanning trees for a given mesh. Their shape can be improved by modifying the Breadth-first-search algorithm so that vertices which are the closest to the singularities or to the boundaries are processed in priority. A spanning tree built this way is presented in red in Fig. 9.

This spanning tree forbids however closed curves around all nodes of the mesh, which is an unnecessarily strong restriction, as we only want to forbid closed curves around the singularities. If we call hanging edge of \mathcal{T} an edge whose leaf node v_l is not a singularity (a leaf node is a node of the graph adjacent to exactly one edge), the branch cut is the sub-graph of \mathcal{T} obtained by successively removing hanging edges until no hanging edges is left in the spanning tree. The result of this subtraction is the branch cut \mathcal{L} needed for the computation of θ . It is depicted in black in Fig. 9.




Figure 9: $\partial\Omega$ is represented in blue, in white the singularities, in red are edges of the spanning tree \mathcal{T} , and in black the edges of the branch cut \mathcal{L} .

5.2.2 Solve the Cauchy-Riemann equations

Once a branch cut \mathcal{L} is available, the field θ can be computed by solving the linear Cauchy-Riemann equations (10). The chosen boundary condition consists in fixing the angle θ at *one* arbitrary point $\mathbf{P} \in \partial\Omega$ so that $\mathcal{C}_{\Omega}(\mathbf{P})$ has one of its branch collinear with $\mathbf{T}(\mathbf{P})$. The problem can be rewritten as the well-posed Eq. (15) and is solved using finite element method on

the triangulation Ω_T with order one Crouzeix-Raviart elements. This kind of elements have proven to be more efficient for cross-field representation as underlined in [7].

$$\begin{cases} \nabla\theta &= \mathbf{n} \times \nabla H \text{ in } \Omega \\ \theta(\mathbf{P}) &= \theta_{\mathbf{P}} \text{ for an arbitrary } \mathbf{P} \in \partial\Omega \\ \theta &\text{ discontinuous on } \mathcal{L} \end{cases} \quad (15)$$

Fig. 10 presents results for 3 different types of singularities. As expected, H and θ isolines are orthogonal, this two scalar fields being harmonic conjugate.

Once H and θ scalar fields are computed on Ω , the cross-field \mathcal{C}_Ω can be retrieved for all $\mathbf{X} \in \Omega$:

$$\mathbf{c}(\mathbf{X}) = \{\mathbf{u}_k = e^{H(\mathbf{X})+jk\theta(\mathbf{X})}, k \in \llbracket 1, 4 \rrbracket\} \quad (16)$$

5.2.3 Non quad-meshable cross-fields

We showed in the previous section that if the imposed singularity pattern respects the condition on Euler's characteristic, then there exists a unique scalar function H (up to an additive constant) verifying Eq. (12) and a unique scalar function θ verifying Eq. (15). This is leading us to a unique cross-field solution \mathcal{C}_Ω to our problem.

However, we never showed that the obtained cross-field \mathcal{C}_Ω is actually respecting the following and mandatory condition:

- $\forall \mathbf{X} \in \partial\Omega$, at least one branch of $\mathcal{C}_\Omega(\mathbf{X})$ is tangent to $\partial\Omega$

Indeed, problem (15) is well-posed and only ensures this condition for a given $\mathbf{P} \in \partial\Omega$. Adding extra boundary conditions for θ will lead to make problem (15) ill-posed. Therefore, when domain Ω is not simply connected, it is possible that obtained a cross-field \mathcal{C}_Ω not tangent on parts of $\partial\Omega$.

When this happens, it means that there is no holomorphic function \mathcal{F} (eq. 6) matching all conditions imposed (*i.e.* singularity pattern and tangency to domain boundaries). Figure 11 shows two different singularity patterns where only the singularity location differs for a same domain Ω . The domain Ω considered is an annulus, and a set of 4 index -1 and 4 index 1 singularities is imposed. Singularity pattern for the example on the left leads to a quad decomposition of the domain and singularity pattern for the example on the right does not. Cross-field obtained for example on the right does not respect the tangency condition on Ω interior boundary. There exists no holomorphic

transformation matching at the same time singularity pattern and cross-field tangency to Ω boundaries.

For the issue of imposing correct singularity configuration we refer the user to Abel-Jacobi conditions given in [24].




Figure 11: Highlighting of singularity patterns for which no holomorphic transformation matching all imposed conditions exists.

5.3 Computing the per-partition parametrization

Obtaining a global parametrization comes down to computing real values functions $U(x, y)$ and $V(x, y)$ defined in Eq. (6), from the known smooth jacobian of the mapping $\mathcal{F} : J_{\mathcal{F}}(\mathbf{z}) = (\partial_U \mathcal{F}(\mathbf{z}), \partial_V \mathcal{F}(\mathbf{z})) \equiv (\tilde{\mathbf{u}}, \tilde{\mathbf{v}})$ Eq. (7). Establishing U and V is equivalent to determining \mathcal{F}^{-1} . We know that $\forall \mathbf{P} = \mathcal{F}(\mathbf{z}) \in \Omega \setminus \mathcal{S}$, $J_{\mathcal{F}}$ corresponds to two branches of the cross-field.

We also know that $J_{\mathcal{F}^{-1}}$ jacobian of \mathcal{F}^{-1} is, $\forall \mathbf{P} \in \Omega \setminus \mathcal{S}$, such as:

$$\begin{cases} J_{\mathcal{F}^{-1}}(\mathbf{P}) &= (\partial_x \mathcal{F}^{-1}(\mathbf{P}), \partial_y \mathcal{F}^{-1}(\mathbf{P})) \\ &= J^{-1}(\mathcal{F}^{-1}(\mathbf{P})) \\ &\equiv (\tilde{\mathbf{u}}, \tilde{\mathbf{v}})^T. \end{cases} \quad (17)$$

which, by using the Eq. (7), gives:

$$\begin{cases} \tilde{\mathbf{u}} &= e^{-H} \mathbf{u} = \nabla U \\ \tilde{\mathbf{v}} &= e^{-H} \mathbf{v} = \nabla V. \end{cases} \quad (18)$$

Now, the parametrization can be simply obtained by solving:

$$\begin{cases} \nabla U = \tilde{\mathbf{u}} \\ \nabla V = \tilde{\mathbf{v}}, \end{cases} \quad (19)$$

using a finite element method with order one Lagrange elements.

We mentioned earlier that for the problem to be well posed, $J_{\mathcal{F}}$ has to be smooth, meaning that $\tilde{\mathbf{u}}$ and $\tilde{\mathbf{v}}$ have to be smooth. This implies that a lifting of the cross-field \mathcal{C}_Ω has to be done on Ω allowing discontinuities of $(\tilde{\mathbf{u}}, \tilde{\mathbf{v}})$ across \mathcal{L} . Once this operation is done, it is possible to determine $(\tilde{\mathbf{u}}, \tilde{\mathbf{v}})$ and compute U and




Figure 10: H isolvalues (in red) and θ isovalues (in blue) for different singularities configurations. The black line represent the cutgraph. All cases have valence 3 singularities imposed at corners (for a total of 4 singularities in each case). The extra singularities which are added inside the domain are: 4 singularities of valence 5 (left figure), 2 singularities of valence 6 (center figure), 1 singularity of valence 8 (right figure).

V by solving the Eq. (19). By construction, U and V will be discontinuous across \mathcal{L} , and their isolines will be in all points tangent to the cross-field \mathcal{C}_Ω .

It is important to note that for this global parametrization computation, no conditions are imposed on the cut graph \mathcal{L} . Therefore, nothing guarantees correspondence of U and V isolines across \mathcal{L} . In principle, this global parametrization is not enough by itself to generate a conforming quadrilateral mesh of a domain Ω . Some works bypass these difficulties [10, 13, 41] by constraining parametric coordinates on boundaries and on \mathcal{L} , however it is not the direction followed in our approach.

The goal here is to obtain per-partition parametrization, starting from the construction of a conformal quad layout of Ω (section 6) and followed by the computation of independent parametrizations for each of the resulting quadrilateral partitions. The per-partition parametrization computed in this manner is aligned with smooth cross-field (singularities can only be located on corners of the partitions). To achieve this goal, the first step is to extract corresponding triangles of each patch. Due to the mutual separatrices' points, neighboring partitions contain the overlapping triangles, as shown in Fig. 13.

An important remark is that no cut graph is needed here to perform a lifting of the underlying cross-field to obtain smooth couple of vector fields $(\tilde{\mathbf{u}}, \tilde{\mathbf{v}})$. As no cut graph is used, all isolines of U and V will be continuous inside the patch, Fig 14.

Based on [7] for locally integrable cross-fields, note

that on each partition holds:

$$\begin{cases} \nabla \times e^{-H}u = 0 \\ \nabla \times e^{-H}v = 0, \end{cases} \quad (20)$$

where from (\tilde{u}, \tilde{v}) is a locally curl-free cross-field. Thus, the resulting parametrization (U, V) is *conformal* on each patch ([11]). Following [42], a *conformal* mapping is also *harmonic*. Radó-Kneser-Choquet (RKC) theorem for planar harmonic mappings (details in [42, 43]) states:

If $f : \mathcal{S} \rightarrow \mathbb{R}^2$ is harmonic and maps the boundary $\partial\mathcal{S}$ homeomorphically into the boundary $\partial\mathcal{Q}$ of some convex region $\mathcal{Q} \subset \mathbb{R}^2$, then f is *one-to-one*.

In our case, the parametric (image) region is in the shape of a rectangle, thus, our parametrization (U, V) is per-partition *bijective*.

6. QUAD LAYOUT: GENERATING AND CORRECTING THE PARTITIONS

In order to reach the ultimate goal of generating a block-structured quad mesh, we must ensure that the quad layout contains only conformal quadrilateral partitions. To do so, a special attention is paid on obtaining the minimal number of separatrices, existence of limit cycles and non-quadrilateral partitions. In the following, we present ways to modify all ill-defined partitions, such that the overall quad layout topology is not jeopardized.

6.1 Obtaining the partitions

To obtain the quad layout, the separatrix-tracing algorithm is performed on retrieved cross-field from H




Figure 12: U isolines (in red) and V isolines (in blue) for different singularities configurations. The black line represents the cutgraph.




Figure 13: Extracted triangulated partition for parametrization

function. Initiation of separatrices in singular areas follows the method from [7], while the tracing of sepa-




Figure 14: U isolines (in red) and V isolines (in blue) on extracted partition

traces in the smooth parts of domain relies on Heun's (a variation of Runge-Kutta 2) numerical scheme.

Our algorithm doesn't have an order of tracing the separatrices - they are all generated simultaneously and independently, similarly to *motorcycle graph* algorithm. Contrary to *motorcycle graph*'s stopping criteria, in our case, the tracing of separatrices is finished as soon as they reach the boundary or encounter a singular point. The number of generated separatrices in this manner is not appropriate for a quad layout, due to the separatrices' initiation process which allows tracing the same separatrix from two different singular points, as shown in Fig. 15. To obtain the valid quad layout, criteria for gaining the minimal number of separatrices are adopted from [7].



Figure 15: Illustration of doubled separatrix and removal of the duplicate

The next step is to establish the existence of limit cycle(s) in the set of the generated separatrices. In more details, we define a separatrix which is transpassing the other separatrices more than once as a *possible* limit cycle. An *authentic* limit cycle is a possible limit cycle which has more far-flung intersections than its pair, or a separatrix not reaching the boundary orthogonally. An authentic limit cycle is then cut at the closest orthogonal intersection with another separatrix, as shown in Fig. 16. By doing so, simplifying the quad layout is accomplished (due to the disappearance of the thin chords), however the obtained partitions can contain T-junctions, which will be modified later on.




Figure 16: Illustration of an authentic limit cycle treatment

6.2 Valid quad layout

If the quad layout contains T-junctions (generated from cutting the authentic limit cycles) or non-quadrilateral partitions, its no longer suitable for block-structured quad meshing. More specifically, a high-quality quad mesh, according to [3], must obtain not only the adequate configuration of singularities (in terms of their number and placement) but also respect the connectivity constraint of being conforming and purely quadrilateral, which the previous issues hinder.

The appearance of T-junctions in our algorithm has already been elucidated, which is not the case with the non-quadrilateral blocks though. In order to examine (and later on correct) the latter case, it's important to mention that our algorithm deals with planar geometries on which, as stressed out by [3] and [14], degenerate singularities (of valence 1 - *singlet* or of valence 2 - *doublet*) may appear. For instance, the typical placement of singularities of valence 2 is on corners of convex boundaries ([3]) or on very sharp/obtuse angles due to the number of cross-field's rotations around the local normal, as the exact formulas for singularity index in [6] show. While these types of singularities do not endanger topological validity of the obtain quad layout, they can cause problems for the block-structured quad meshing by forming non-quadrilateral blocks, as shown in Fig. 18.

6.3 Correcting the partitions

When the partition contains a T-junction, the idea is to find the closest neighbor (which is not the boundary vertex or another T-junction) and, by knowing the parametrization of these blocks, to merge these two separatrices in the parametric space and correct all patches affected by this action, Fig 17. This process is iteratively repeated until the quad layout is T-junction free. Following the definition of separatrices and singular points, each of the newly obtained partitions has smooth cross-field inside allowing further on parameterization.




Figure 17: Modifying patches containing T-junctions and affected partitions

Correcting the triangular partitions is performed by forcing the singularity of valence 2 on the boundary into splitting in two singularities of valence 3, so that the topology of the overall obtained partitions remains valid (Fig. 18). The position of the imposed singularity is determined simply by finding the barycenter of the partition. Final quad layout is obtained by splitting the triangular patch into three quadrilateral partitions and correcting all patches influenced by the former action.




Figure 18: Modifying the patch containing a singularity of valence 2 and affected partitions

7. QUAD MESHING

To generate a final quad mesh of the model, we start by remeshing each layout's partition with bilinear transfinite interpolation (TFI) in the parametric space, while respecting the discretizations of layout's edges implied by the field H . This step is followed by a computation of a per-partition parametrization aligned with the cross-field (as detailed in section 5.3). Finally, generating a quad mesh of pleasing quality and intact singularities' positions (examples in section 8) is obtained by mapping TFI mesh onto the physical space. At this point, it is indeed possible to further optimize vertices' placements via smoothing (e.g. Winslow from [44, 45]), in the cases when the application allows/encourages moving singularities away from their initial positions. For the sake of completeness, details of the above-mentioned steps are elucidated in the following.

In order to remesh each patch of the quad layout, the chords' data has been extracted. Keeping in mind that the quad layout is now composed of only conforming quadrilaterals, it is possible to consistently discretize all its edges by following the size map implied by H :

$$s = e^H, \quad (21)$$

where s represents the mesh size. This step tends to generate a uniform quad mesh, unless a specific size map (via implying an adequate configuration of singularities) is imposed by the user. With obtained discretization of edges, a quad mesh in parametric space is trivially created by remeshing each partition with bilinear transfinite interpolation ([46]).

As the parametrization is now defined, the mapping of the quad mesh from parametric onto physical space is achieved by simply finding the triangle T_M in the parametric space to whom the point $M(U, V)$ belongs. Afterwards, the computing of coordinates $M(x, y, z)$ is done by assuming that there exists a *linear mapping* from physical space X to parametric space P inside the triangle T_M .

8. CONCLUSION AND FUTURE WORK

We presented an algorithm that can effectively generate a conformal quad layout from a naturally appearing as well as user-imposed set of singularities, as shown in Fig. 19. The latter case is possible under constraints of the topological correctness and existence of conformal mapping for imposed singularity pattern. In particular, the developed method relies on the computation of two cross-fields. The first one is computed solving a non-linear system and is used to locate singularities and obtain their valence, the second one is computed on a refined triangular mesh to accommodate high gradients and improve cross-field representation around singularities. It is obtained solving two linear systems. With the correction scheme for limit cycles and non-quadrilateral partitions, we can produce a conformal quad layout and, finally, a block-structured quad mesh. Remeshing of each patch relies on constructing a per-partition parametrization and following the isolines from the second cross-field. A quad mesh obtained in this manner preserves singularity placement and is of pleasing quality, as shown in Table 1. Once the quad layout obtained, the computational cost of our quad meshing algorithm depends on the target number of quads or the approximate length of the edge imposed by the user, which takes less than a second in the presented examples. At last, our novel pipeline exhibits interactive design while being simple, automatic and available in the Gmsh software.

Attractive directions for future work include, although are not limited to extending the algorithm for computations on various surfaces, which is currently underway, as well as its adaptation for working on varying user-prescribed element sizes. In the first case, the extension of the algorithm includes a different computational procedure for the H function and an updated

scheme for dealing with limit cycles and T-junctions. If the anisotropic quad mesh is required, two scalings for the cross-field must be defined instead of only one like in our case.

Geometry	Edge length	Mesh quality		
		$\bar{\eta}$	η_ω	τ
Figure 3	0.05	0.96	0.74	88.16
Figure 4	0.02	0.98	0.40	98.39
Model 1	0.05	0.96	0.34	93.38
Model 2	0.05	0.94	0.47	89.44
Model 3	0.05	0.94	0.52	89.19
Model 4	0.20	0.98	0.76	96.93
Model 5	0.20	0.98	0.63	96.94
Model 6	0.05	0.88	0.70	48.91
Model 7	0.02	0.98	0.78	98.78
Model 8	0.02	0.97	0.51	98.28
Model 9	0.02	0.98	0.74	96.36
Model 10	0.02	0.97	0.37	95.03

Table 1: The quality of obtained meshes: average quality of elements under $\bar{\eta}$, the worst element’s quality under η_ω and percentage of elements with quality higher than 0.9 under τ . Element’s quality is adopted from [47], and is a measure of the deviation of element’s angle from $\frac{\pi}{2}$.

9. ACKNOWLEDGMENTS

The present study was carried out in the framework of the research project ”Hextreme”, funded by the European Research Council (ERC-2015-AdG-694020) and hosted at the Université catholique de Louvain.

References

- [1] Campen M. “Partitioning surfaces into quadrilateral patches: a survey.” *Computer Graphics Forum*, vol. 36, pp. 567–588. Wiley Online Library, 2017
- [2] Beaufort P.A., Lambrechts J., Henrotte F., Geuzaine C., Remacle J.F. “Computing cross fields A PDE approach based on the Ginzburg-Landau theory.” *Procedia engineering*, vol. 203, 219–231, 2017
- [3] Bommès D., Lévy B., Pietroni N., Puppo E., Silva C.T., Tarini M., Zorin D. “Quad Meshing.” *Eurographics (STARs)*, pp. 159–182. 2012
- [4] Kowalski N., Ledoux F., Frey P. “A PDE based approach to multidomain partitioning and quadrilateral meshing.” pp. 137–154, 2013
- [5] Fogg H.J., Armstrong C.G., Robinson T.T. “Automatic generation of multiblock decompositions

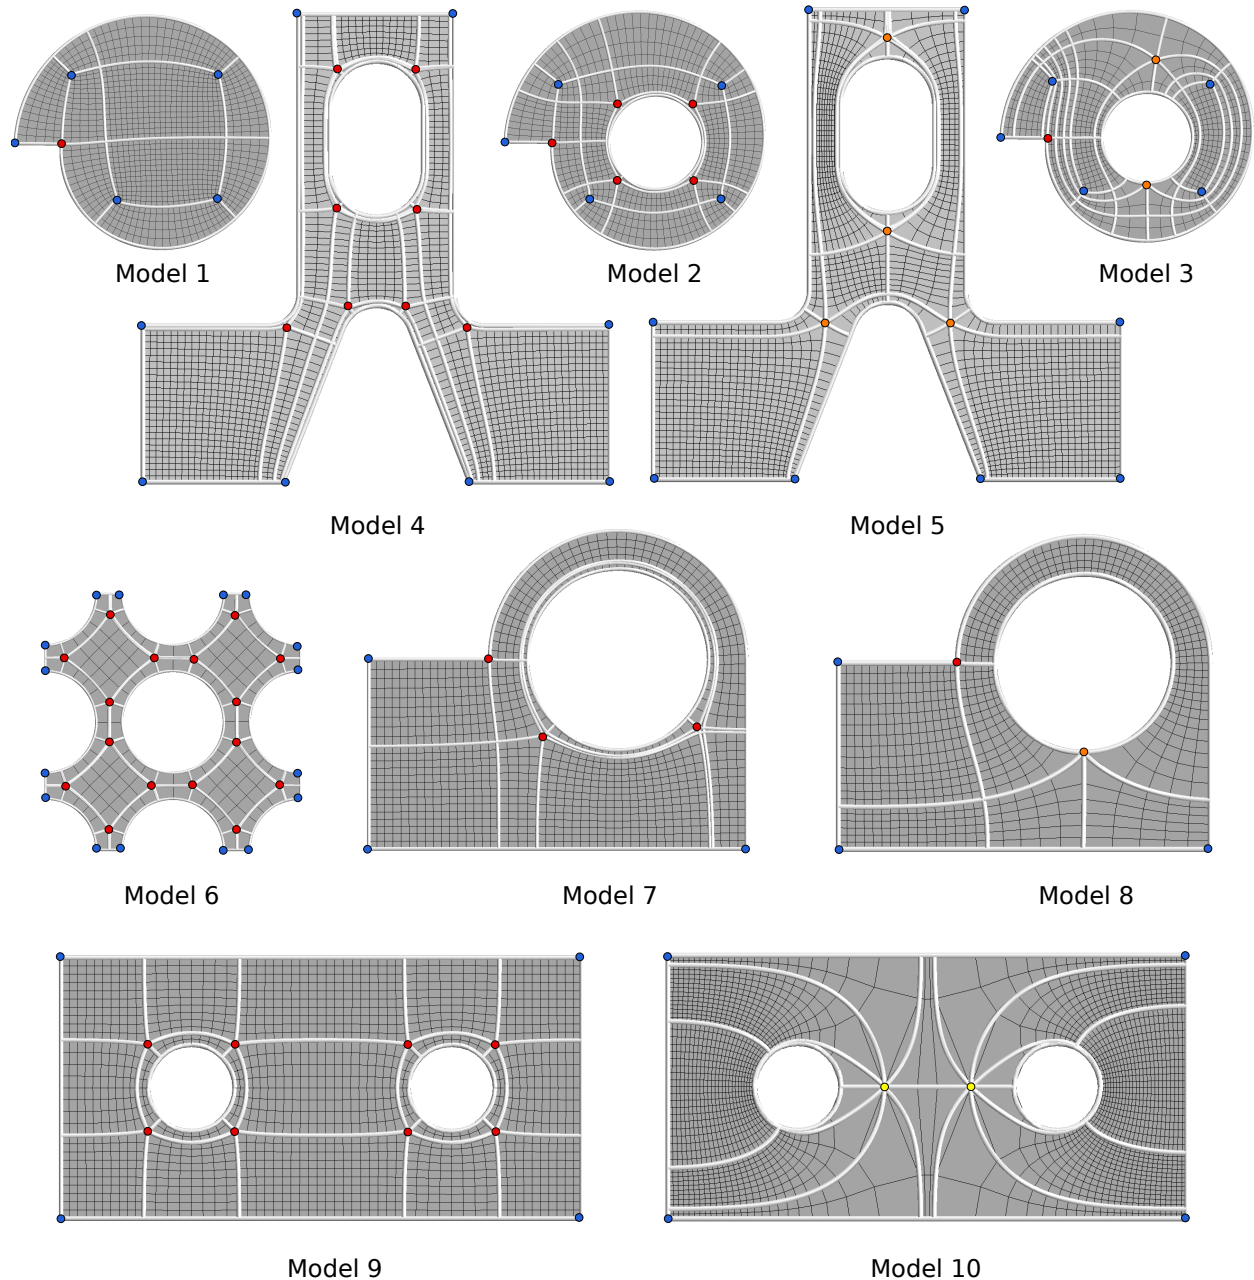


Figure 19: Examples of quad meshes generated from natural and user-imposed singularity configurations. Singularities of valence 3 are marked in blue, of valence 5 in red, of valence 6 in orange and of valence 8 in yellow.

of surfaces.” *International Journal for Numerical Methods in Engineering*, vol. 101, no. 13, 965–991, 2015

[6] Viertel R., Osting B. “An Approach to Quad Meshing Based on Harmonic Cross-Valued Maps and the Ginzburg–Landau Theory.” *SIAM Journal on Scientific Computing*, vol. 41, no. 1, A452–

A479, 2019

[7] Jezdimirovic J., Chemin A., Rémacle J.F. “Multi-block decomposition and meshing of 2D domain using Ginzburg–Landau PDE.” *28th International Meshing Roundtable*, 2019

[8] Campen M., Kobbelt L. “Quad layout embedding via aligned parameterization.” *Computer Graph-*

- ics Forum*, vol. 33, pp. 69–81. Wiley Online Library, 2014
- [9] Campen M., Bommès D., Kobbelt L. “Quantized global parametrization.” *ACM Transactions on Graphics (TOG)*, vol. 34, no. 6, 192, 2015
- [10] Bommès D., Campen M., Ebke H.C., Alliez P., Kobbelt L. “Integer-grid maps for reliable quad meshing.” *ACM Transactions on Graphics (TOG)*, vol. 32, no. 4, 98, 2013
- [11] Ray N., Li W.C., Lévy B., Sheffer A., Alliez P. “Periodic global parameterization.” *ACM Transactions on Graphics (TOG)*, vol. 25, no. 4, 1460–1485, 2006
- [12] Geuzaine C., Remacle J.F. “Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities.” *International journal for numerical methods in engineering*, vol. 79, no. 11, 1309–1331, 2009
- [13] Bommès D., Zimmer H., Kobbelt L. “Mixed-integer quadrangulation.” *ACM Transactions On Graphics (TOG)*, vol. 28, no. 3, 1–10, 2009
- [14] Ray N., Vallet B., Li W.C., Lévy B. “N-symmetry direction field design.” *ACM Transactions on Graphics (TOG)*, vol. 27, no. 2, 1–13, 2008
- [15] Vaxman A., Campen M., Diamanti O., Panozzo D., Bommès D., Hildebrandt K., Ben-Chen M. “Directional field synthesis, design, and processing.” *Computer Graphics Forum*, vol. 35, pp. 545–572. Wiley Online Library, 2016
- [16] Diamanti O., Vaxman A., Panozzo D., Sorkine-Hornung O. “Integrable polyvector fields.” *ACM Transactions on Graphics (TOG)*, vol. 34, no. 4, 38, 2015
- [17] Crane K., Desbrun M., Schröder P. “Trivial connections on discrete surfaces.” *Computer Graphics Forum*, vol. 29, pp. 1525–1533. Wiley Online Library, 2010
- [18] Ray N., Vallet B., Alonso L., Levy B. “Geometry-aware direction field processing.” *ACM Transactions on Graphics (TOG)*, vol. 29, no. 1, 1–11, 2009
- [19] Scheuermann G., Hagen H., Kruger H., Menzel M., Rockwood A. “Visualization of higher order singularities in vector fields.” *Proceedings. Visualization’97 (Cat. No. 97CB36155)*, pp. 67–74. IEEE, 1997
- [20] Scheuermann G., Kruger H., Menzel M., Rockwood A.P. “Visualizing nonlinear vector field topology.” *IEEE Transactions on Visualization and Computer Graphics*, vol. 4, no. 2, 109–116, 1998
- [21] Tricoche X., Scheuermann G., Hagen H. “Higher order singularities in piecewise linear vector fields.” *The Mathematics of Surfaces IX*, pp. 99–113. Springer, 2000
- [22] Günther T., Theisel H. “The state of the art in vortex extraction.” *Computer Graphics Forum*, vol. 37, pp. 149–173. Wiley Online Library, 2018
- [23] Weinkauff T., Theisel H., Shi K., Hege H.C., Seidel H.P. “Extracting higher order critical points and topological simplification of 3D vector fields.” *VIS 05. IEEE Visualization, 2005.*, pp. 559–566. IEEE, 2005
- [24] Lei N., Zheng X., Luo Z., Luo F., Gu X. “Quadrilateral mesh generation II: Meromorphic quartic differentials and Abel–Jacobi condition.” *Computer Methods in Applied Mechanics and Engineering*, vol. 366, 112980, 2020
- [25] Barnette D., Jucovič E., Trenkler M. “Toroidal maps with prescribed types of vertices and faces.” *Mathematika*, vol. 18, no. 1, 82–90, 1971
- [26] Jucovič E., Trenkler M. “A theorem on the structure of cell-decompositions of orientable 2-manifolds.” *Mathematika*, vol. 20, no. 1, 63–82, 1973
- [27] Izmetiev I., Kusner R.B., Rote G., Springborn B., Sullivan J.M. “There is no triangulation of the torus with vertex degrees 5, 6, ..., 6, 7 and related results: Geometric proofs for combinatorial theorems.” *Geometriae Dedicata*, vol. 166, no. 1, 15–29, 2013
- [28] Viertel R., Osting B., Staten M. “Coarse Quad Layouts Through Robust Simplification of Cross Field Separatrix Partitions.” *arXiv preprint arXiv:1905.09097*, 2019
- [29] Myles A., Pietroni N., Zorin D. “Robust field-aligned global parametrization.” *ACM Transactions on Graphics (TOG)*, vol. 33, no. 4, 1–14, 2014
- [30] Ray N., Sokolov D. “Robust polylines tracing for n-symmetry direction field on triangulated surfaces.” *ACM Transactions on Graphics (TOG)*, vol. 33, no. 3, 1–11, 2014
- [31] Jadhav S., Bhatia H., Bremer P.T., Levine J., Nonato L.G., Pascucci V. “Consistent approximation of local flow behavior for 2d vector fields using edge maps.” Tech. rep., Lawrence Livermore National Lab.(LLNL), Livermore, CA (United States), 2011

- [32] Eppstein D., Goodrich M.T., Kim E., Tamstorf R. “Motorcycle graphs: canonical quad mesh partitioning.” *Computer Graphics Forum*, vol. 27, pp. 1477–1486. Wiley Online Library, 2008
- [33] Borden M.J., Benzley S.E., Shepherd J.F. “Hexahedral Sheet Extraction.” *IMR*, pp. 147–152. 2002
- [34] Daniels J., Silva C.T., Shepherd J., Cohen E. “Quadrilateral mesh simplification.” *ACM transactions on graphics (TOG)*, vol. 27, no. 5, 1–9, 2008
- [35] Hertzmann A., Zorin D. “Illustrating smooth surfaces.” *Proceedings of the 27th annual conference on Computer graphics and interactive techniques*, pp. 517–526. 2000
- [36] Palmer D., Bommers D., Solomon J. “Algebraic representations for volumetric frame fields.” *ACM Transactions on Graphics (TOG)*, vol. 39, no. 2, 1–17, 2020
- [37] Knöppel F., Crane K., Pinkall U., Schröder P. “Globally optimal direction fields.” *ACM Transactions on Graphics (ToG)*, vol. 32, no. 4, 1–10, 2013
- [38] Kronecker L. “Über Systeme von Funktionen mehrerer Variablen, Monatsber. d.” *Kgl. Pr. Akad. Berlin*, p. 695, 1869
- [39] Bunin G. “A continuum theory for unstructured mesh generation in two dimensions.” *Computer Aided Geometric Design*, vol. 25, no. 1, 14–40, 2008
- [40] Bethuel F., Brezis H., Hélein F., et al. *Ginzburg-Landau Vortices*, vol. 13. Springer, 1994
- [41] Kälberer F., Nieser M., Polthier K. “Quadcover-surface parameterization using branched coverings.” *Computer graphics forum*, vol. 26, pp. 375–384. Wiley Online Library, 2007
- [42] Floater M.S., Hormann K. “Surface parameterization: a tutorial and survey.” *Advances in multiresolution for geometric modelling*, pp. 157–186. Springer, 2005
- [43] Iwaniec T., Onninen J. “Radó-Kneser-Choquet Theorem for simply connected domains (p -harmonic setting).” *Transactions of the American Mathematical Society*, vol. 371, no. 4, 2307–2341, 2019
- [44] Winslow A.M. “Numerical solution of the quasilinear Poisson equation in a nonuniform triangle mesh.” *Journal of computational physics*, vol. 1, no. 2, 149–172, 1966
- [45] Knupp P.M. “Winslow smoothing on two-dimensional unstructured meshes.” *Engineering with Computers*, vol. 15, no. 3, 263–268, 1999
- [46] Thompson J.F., Soni B.K., Weatherill N.P. *Handbook of grid generation*. CRC press, 1998
- [47] Remacle J.F., Lambrechts J., Seny B., Marchandise E., Johnen A., Geuzainet C. “Blossom-Quad: A non-uniform quadrilateral mesh generator using a minimum-cost perfect-matching algorithm.” *International journal for numerical methods in engineering*, vol. 89, no. 9, 1102–1119, 2012