Quad Layouts via Constrained T-Mesh Quantization

M. Lyon[1] , M. Campen[2] , and L. Kobbelt[1]

[1]RWTH Aachen University, Germany
[2]Osnabrück University, Germany

Figure 1: *Given a seamless parametrization of a surface as input we construct a non-conforming T-mesh aligned with it. A modified version of the motorcycle graph is employed for this purpose, in which traces do not stop at the first collision. We solve an integer linear program to assign integers to the arcs of this T-mesh. By carefully constraining this quantization, the result implies a coarse conforming quad layout whose separatrices deviate less than a user given bound from the directions prescribed by the seamless input parametrization.*

Abstract
We present a robust and fast method for the creation of conforming quad layouts on surfaces. Our algorithm is based on the quantization of a T-mesh, i.e. an assignment of integer lengths to the sides of a non-conforming rectangular partition of the surface. This representation has the benefit of being able to encode an infinite number of layout connectivity options in a finite manner, which guarantees that a valid layout can always be found. We carefully construct the T-mesh from a given seamless parametrization such that the algorithm can provide guarantees on the results' quality. In particular, the user can specify a bound on the angular deviation of layout edges from prescribed directions. We solve an integer linear program (ILP) to find a coarse quad layout adhering to that maximal deviation. Our algorithm is guaranteed to yield a conforming quad layout free of T-junctions together with bounded angle distortion. Our results show that the presented method is fast, reliable, and achieves high quality layouts.

CCS Concepts
• Computing methodologies → Computer graphics; Mesh models; Mesh geometry models; Shape modeling;

1. Introduction

Quad meshes are the preferred discrete surface representation for many shape modeling applications of design and engineering alike. Therefore, the automatic generation of such meshes has been an ongoing topic of research in computer graphics. For a high quality quad mesh, individual elements should have angles close to $\pi/2$ and be aligned in certain ways, e.g. to the underlying surface's principal curvature directions. A variety of approaches have been explored for the generation of quad meshes [BLP*13]. A class of algorithms with a particular focus on element shape and alignment quality is that of parametrization-based field guided methods [KNP07, BZK09, KMZ11, BCE*13, PPTSH14, ESCK16].

An additional quality criterion that these methods, however, do not explicitly promote is the global structure of the mesh, in particular the simplicity of the mesh's so-called *block structure* or *base complex* [BLP*13, §1.1]. This aspect is closely related to the question how the irregular vertices are connected in the mesh by sequences of edges. This connectivity constitutes the mesh's *quad layout* [Cam17]. If the quad mesh has a simple, i.e., a coarse layout it can be viewed as a "mesh of meshes", a coarse quad partition with finer regular quad grids inside each patch (Figure 1 right). This enables the construction of mesh hierarchies, the structured parametrization of the mesh over simple domains, or the definition of spline spaces on top of the mesh [TPP*11, MAB*19, HSJ*20].

In this paper we present an algorithm for the creation of coarse quad layouts on 2-manifold surfaces. Such layouts can then, for instance, be refined to block-structured quad meshes, or be passed as

connectivity constraints to the above mentioned (natively layout-unaware) quad mesh generation techniques. The layout construction – akin to the above parametrization-based mesh generation methods – is built on top of a seamless surface parametrization. This allows us to offer explicit control over the layout's singularity configuration and its directional alignment, to reliably yield valid layouts, and to share a common foundation with mesh generation techniques for seamless integration. As a particular feature, our method offers explicit and precise control over the balance between the two key opposing objectives inherent to quad layouts: coarseness and directional alignment.

Problem Statement: Given a seamless surface parametrization with arbitrary singularities on a surface of arbitrary topology, generate a coarse quad layout with exactly these singularities as irregular nodes such that its arcs (also called separatrices) do not directionally deviate from the parametrization's isolines by more than a user-given angular bound α.

After reviewing previous methods for the creation of coarse quad layouts in Section 2, we describe in Section 3 how we construct a T-mesh by tracing parametric iso-lines of a seamless parametrization taken as input. In our method we implicitly encode layout connectivity via a discrete function on this T-mesh. Section 4 details this encoding and shows how non-negative integers can be assigned to the T-mesh's arcs such that a valid and high quality layout is implied. Such an assignment, called quantization, can be found by solving an integer linear program as presented in Section 5, before ultimately making the layout explicit in Section 6. We show the effectiveness of our algorithm on a variety of examples in Section 7. Figure 1 illustrates the process.

2. Related Work

Early work that involved the generation of quad layouts as a sub-step produced rather unstructured layouts without any particular form of shape-aware directional alignment [EH96, BMRJ04]. Later this problem received dedicated attention and subsequent methods often take some form of directional guidance into account, whether by means of the underlying surface's principal curvature directions, or more flexibly and controllably by specifically designed or prescribed cross fields, quad meshes, or seamless parametrizations.

Layout Simplification. For instance, the method by Bommes et al. [BLK11] takes as input a quad mesh with a possibly dense base complex and iteratively modifies it so as to remove certain helical connectivity patterns – which are one, though not the sole cause of low quality base complexes. A coarser quad layout can then be extracted from the modified quad mesh.

Tarini et al. [TPP*11] follow a similar strategy but enable more general modifications by working directly on separatrices (i.e., the paths forming the layout's arcs, its patches' borders). They iteratively improve a layout energy which is based on length and direction deviation of the separatrices by removing a separatrix and bringing the then incomplete layout back into a (coarser) complete state by a series of separatrix reconnections and an insertion.

Instead of starting with conforming dense quad layouts, Viertel et al. [VOS19] start from an initial non-conforming layout, a T-mesh with many T-junctions. It is obtained by tracing streamlines of a surface cross field. In this layout so-called chord collapses are applied greedily, in order from narrow to wide, while excluding collapses that would result in too much directional deviation. Particularly on complex or closed surfaces, T-junctions may remain in the final layout, making it non-conforming.

These methods have in common that modifications are applied iteratively in a greedy fashion.

Layout from Separatrix Candidates. A common strategy to create layouts from scratch is based on finding a set of separatrix candidates, i.e., paths connecting pairs of prescribed (irregular) layout nodes in topologically distinct ways, from which a subset is then selected to define a complete layout.

Razafindrazaka et al. [RRP15] trace isolines of a seamless parametrization, starting from its singularities (which form the layout's irregular nodes). Whenever two traces meet, this implies a separatrix candidate between their two origin singularities. Each candidate is associated with a cost, penalizing directional deviation from the parametric isolines. A binary problem is then solved to select a cost-minimizing subset that properly connects all singularities without crossing in improper ways. In theory the candidate set is infinite; in practice one needs to restrict to a finite subset. This may preclude the existence of a valid solution (or a high-quality solution). Tracing up to a maximum distance limit is reported to commonly work well, but the existence of a solution is not guaranteed (unless trial-and-error with increasing distance limit is performed).

Pietroni et al. [PPM*16] follow a similar approach but create separatrix candidates based on a cross field [VCD*16] rather than a parametrization (which is harder to obtain with the required properties [CSZZ19]). Similar to [RRP15] they solve a binary linear program to choose a non-conflicting subset of these. As a consequence of not deriving directional guidance from a parametrization (which corresponds to an integrable cross field) but rather from a generic cross field, a complete conforming quad layout cannot be guaranteed; T-junctions have to be accepted, similar to the approach of Viertel et al. [VOS19] discussed above. T-junctions may be reduced by increasing the number of separatrix candidates established for each singularity or by iterating the process with fixed partial layout, but complete removal can only be achieved by inserting additional irregularities.

For both these methods, runtime, quality, and even the existence of a solution depend on the precomputed set of separatrix candidates. If the set is chosen too small there may be no valid subset. Increasing the candidate set size at the cost of increased runtime only increases the probability of program feasibility. In contrast, our formulation, instead of using a binary program which picks separatrices from a finite set of candidates, employs an integer program which can choose from an infinite set of separatrices and is always guaranteed to be feasible.

[RP17] and [ZZY16] propose further candidate set based methods, but start from a quad mesh as input. This effectively enables a fallback to the quad mesh itself or its base complex as a valid (though commonly rather dense) output layout in case no other solution is found due to the involved restriction.

Dual Approach. Campen et al. [CBK12] tackle the problem in the dual space and search for a set of dual loops which separate all singularities. While the dual setting allows for relatively simple conditions guaranteeing a valid primal layout it does not allow explicit control over the geometric quality of the implied separatrices.

Mesh as Layout. An alternative approach to create coarse quad layouts is to use a generic quad mesh generation method (without explicit layout considerations) and aim for a mesh with very large elements – which can then be considered a coarse layout. While some of these methods are robust enough to operate reliably under this extreme requirement (e.g. [BCE*13, CBK15]), it is typically difficult for these methods to control the quality of the resulting layout under these circumstances (cf. Section 7.1).

3. T-Mesh Construction

A motorcycle graph is a cell decomposition of a surface [EGKT08]. Campen et al. [CBK15, MC19] show how a rectangular partition of a surface can be constructed by tracing motorcycles starting at singularities along the iso-lines of a seamless parametrization. The resulting T-mesh $\mathcal{T} = (\mathcal{N}, \mathcal{A}, \mathcal{P})$ consists of nodes $\mathcal{N}$ for every singularity and intersection of traces, arcs $\mathcal{A}$ consisting of the segments of a trace between two nodes, and patches $\mathcal{P}$ representing rectangular regions bounded by arcs.

Such a motorcycle graph forms the basis of our quantization (cf. Section 4) which assigns integer lengths to the arcs and thereby defines the connectivty of the resulting layout.

We adopt a construction similar to [CBK15], with one important difference: in our case motorcycles may survive a crash with an existing trace and continue driving. In such a case a regular valence 4 node is formed at the intersection. Before defining the criterion that determines this, we establish a few definitions and notation.

Notation. Given two traces t_i and t_j starting in singularities i and j, respectively, we refer to the node created at their intersection as $n_{ij} \in \mathcal{N}$. (For simplicity of notation we ignore the fact that two traces may intersect more than once; it will be clear from the context which intersection node is referred to.) We define $S_{ij} \subset \mathcal{A}$ as the set of arcs between the start of trace t_i and node n_{ij}, and $l_{ij} \in \mathbb{R}$ as the total parametric length of those arcs. The two arc sets S_{ij} and S_{ji} form the legs of a right triangle. Let $\alpha_{ij} \in (-\frac{\pi}{2}, \frac{\pi}{2})$ be the signed (ccw) angle of that triangle at the start of trace t_i.

We further give a dedicated name, n_{i*}, per trace t_i to the intersection closest to the origin of t_i which satisfies $l_{i*} > l_{*i}$, i.e. the first intersection of t_i with a trace that starts inside the $\pi/2$-sector (blue) around t_i (see inset). In the classical motorcycle graph a trace t_i would stop exactly at n_{i*}. By contrast, we keep on tracing and only stop based on the following criterion.

Stopping Criterion. Given a user defined angular bound α, a trace is stopped as soon as it intersected two traces t_k and t_l such that


Figure 2: *T-meshes on the* BUNNY *model created for angle deviation bound $\alpha = 5°, 15°, 35°$, respectively.*

$\alpha_{ik} \in [0, \alpha]$ and $\alpha_{il} \in [-\alpha, 0]$, i.e., as soon as on both sides of trace t_i a trace is found such that the formed triangles are contained in the sector of half-angle α around t_i. With this stopping criterion we ensure that we reach the traces of at least one singularity on each side of the trace which, intuitively, could be connected to singularity i by a separatrix which respects the maximum angular deviation bound. We will see in detail in Section 4.3 how this construction helps to guarantee a maximum separatrix deviation. Figure 2 shows a few examples of the resulting T-meshes for different α and how tighter bounds lead to longer traces.

4. Quantization

A quantization of a T-mesh is an assignment of a non-negative integer q_i to every arc $a_i \in \mathcal{A}$. These values are to be interpreted as parametric length specification for the arcs, virtually overriding their length in the input parametrization. Since every arc is associated with a parametric iso-line, the quantization implies an assignment of integer parametric distances between singularities [CBK15]. These distances can be used as constraints for a global re-parametrization of the surface where every singularity is located on an integer position, enabling the extraction of a quad mesh [CBK15, LCBK19]. This quad mesh's base complex defines a quad layout. The edges of the base complex – the separatrices – connect singularities that lie on the same parametric iso-line. Effectively, the quantization implies which singularities are connected by separatrices by assigning zero lengths to certain arcs. Depending on the T-mesh structure and the quantization, singularities may be connected by separatrices corresponding to different parametric directions (cf. Figure 3(left, middle)) or not (cf. Figure 3(right)).

While the above previous works aim to find a quantization which approximates the arcs' original parametric lengths, we detail in this section how we instead find a quantization that promotes a coarse quad layout. We begin by discussing two important properties of a quantization – *validity* and *consistency* – in the following section. We continue by presenting a sufficient condition which guarantees validity, and finally describe which additional constraints we set up to enforce a high quality quad layout.

4.1. Consistency and Validity

In order to compute a valid (locally injective, fold-over free) parametrization that adheres to the singularity distance constraints defined by the quantization, two properties need to be fulfilled. The quantization needs to be consistent and valid [CBK15].

Figure 3: *Layouts resulting from different quantizations. Singularities (dots) are separated in vertical direction leading to an angle $\alpha > \pi/4$ compared to the desired direction (left). With our separation constraint, singularities are separated in the direction corresponding to the larger difference in the seamless parametrization always leading to separatrix angles smaller than $\pi/4$ (middle). If both separatrices form an unacceptably large angle, singularities need to be separated in both directions (right). Note that the actual quantization (beyond zero vs non-zero) is less important for the layout, consisting of the (bold) edges of the base complex only.*

A quantization is *consistent* if pairs of opposite sides of each T-mesh patch are quantized to the same length. This property ensures that all patches remain rectangular in the parametric domain.

A quantization is *valid* if the distance between any two singularities is strictly positive, such that they do not collapse parametrically. In terms of the quantization this means there must be no arc path with total quantized length of 0 between any two singularities.

Let $\mathbf{d}_{ij} \in \mathbb{R}^2$ be the (u, v)-difference between singularities i and j measured in the seamless parametrization (in some path homotopy class) and $\mathbf{q}_{ij} \in \mathbb{Z}^2$ be the quantized difference vector. W.l.o.g. assume that $\mathbf{d}_{ij}^u > 0$ and $|\mathbf{d}_{ij}^u| \geq |\mathbf{d}_{ij}^v|$. To ensure validity it is sufficient to require $\mathbf{q}_{ij}^u > 0 \vee \mathbf{q}_{ij}^v \neq 0$ [CBK15]. We propose to use instead the simpler sufficient condition of $\mathbf{q}_{ij}^u > 0$, i.e. singularities are separated by the quantization in the dominant separation direction in the input parametrization. While this condition is stricter, what we effectively exclude are quantizations with $\mathbf{q}_{ij}^u = 0$ and $\mathbf{q}_{ij}^v \neq 0$ (Figure 3 left); thereby we already ensure a maximal separatrix deviation of $\pi/4$. Even more practically relevant, this disjunction-less condition is easier to formulate and more efficient to handle.

Campen et al. propose for their iterative quantization improvement to test validity after each change by explicitly searching singularity connecting arc paths with quantized length of zero. In the worst case this test is in $O(n^2)$ and in general not easily applicable in the context of solving an integer linear program. In the following section we present how separation of singularities can be guaranteed using only one linear constraint per trace.

4.2. Singularity Separation

To guarantee the separation of singularities in the quantization we propose a simple criterion based on the following lemma:

Lemma 1. *If for every trace t_i one of the arcs in S_{i*} is quantized to at least 1, the quantization is valid.*

Here, S_{i*} is the set of arcs between the start of t_i and n_{i*}, i.e. the first intersection of t_i with an earlier trace (cf. Section 3).

Proof. We show that every singularity in the $\pi/2$-sector centered at a singularity i and around the positive u direction is separated from i. The same argument then trivially holds for all other parametric directions as well.

Let j be the singularity at the start of t_*. Then, j is separated from i in u direction by at least 1 according to the lemma's premise. Consider another singularity k within the $\pi/2$-sector and its trace t_k in negative (or positive) v direction towards t_i. If t_k intersects the parametric iso-line of t_i it must do so further away than t_* since t_* is the intersection closest to i by definition. Since the intersection lies behind that of t_*, k is separated by at least 1 in u direction as well. If, on the other hand, the trace of k does not intersect the parametric iso-line of t_i it must have been stopped before reaching it. Since at least one of the arcs of t_k is quantized to at least one according to the condition in the lemma, k is separated from i in positive (or negative) v direction. $\square$

Note that Lemma 1 does not depend on our special T-mesh construction and also holds if traces stop at the first intersection.

We conclude that it is sufficient to require one arc of S_{i*} to be quantized to a strictly positive number for each trace t_i to ensure validity, i.e., one constraint per trace is sufficient.

4.3. Layout Constraints

In the previous section we detailed how simple constraints on the minimal quantization of certain arcs guarantee a valid quantization. Computing the smallest quantization adhering to these constraints results in a coarse quad layout. However, the separatrices of that layout may deviate up to $\pi/4$ from the directions of the seamless parametrization (cf. Figure 3 b). In particular the separatrix corresponding to a trace t_i will connect to the first singularity within the $\pi/2$-sector which is not separated in the direction orthogonal to t_i. We propose a simple solution to create layouts in which the separatrices do not deviate from the intended directions more than a user specified maximum of α by enforcing the additional separation of offending singularities in orthogonal direction (cf. Figure 3 c).

Given two intersecting traces t_i and t_j with $l_{ij} \geq l_{ji}$ originating from singularites i and j, we defined in the previous sections constraints that ensure the arc set S_{ij} is quantized to at least 1, separating i and j in the direction of t_i. If S_{ji} will be quantized to 0, i and j would lie on the same iso-parameter line and would therefore be connected by a separatrix (unless there is a singularity closer to i which also lies on the same iso-parameter line). If angle $|\alpha_{ij}|$ of this separatrix is larger than the user defined bound α we need to prevent such a quantization. By additionally separating j in the direction of t_j it is ensured that the corresponding separatrix is not formed. We therefore include the additional constraint that S_{ji} is quantized to at least 1 for every pair of intersecting traces t_i and t_j with $|\alpha_{ij}| > \alpha$.

In the light of this, we can now explain the rationale behind our choice of stopping criterion in the T-mesh construction, continuing a trace t_i until at least two traces t_k and t_l are intersected

with $\alpha_{ik} \in [0, \alpha]$ and $\alpha_{il} \in [-\alpha, 0]$ (cf. Section 3). We want to guarantee that all separatrix deviations are below the user defined bound. With the constraints above, any separatrix created between singularities whose traces intersect fulfills this. But what if the quantization implies none of them, i.e. S_{ji} is quantized to at least 1 for all traces t_j that intersect t_i, which may be necessary due to a combination of separation and consistency constraints? In that case the separatrix will lie in the corridor (blue) between the two potential separatrices connecting the start of t_i to the starts of t_k or t_l. Since both these potential separatrices satisfy the bound, the actually implied separatrix (blue dotted line) lying in between satisfies it as well.

4.4. Feature Lines & Boundaries

Some models contain sharp creases for which it is typically desirable that they are represented by arcs in a layout. Similarly, boundaries of models should be represented by arcs. Both these cases, as well as arbitrary, user defined feature curves are supported by our method by aligning the parametrization with these features and tracing them with motorcycles to integrate them into our T-mesh [CBK15]. To ensure that the traced separatrix is not diverted towards another singularity away from the feature curve it is enough to simply add layout constraints (cf. Section 4.3) for *all* intersecting traces. This can be interpreted as prescribing a maximum separatrix deviation of $\alpha = 0$ for all traces that follow a feature curve. Figure 9 shows examples where boundary alignment is enforced this way.

5. Integer Linear Program

We established how a quantization can be constrained such that no separatrix in the implied quad layout deviates more than a user given bound from the intended direction. In the space of feasible quantizations respecting these constraints, we are looking for a quad layout as coarse as possible. In this section we discuss how such a quantization can be found efficiently by solving an integer linear program (ILP). We first describe the basic integer linear program, which can be constructed in a straightforward fashion following the previous discussion. After that we discuss how the program size can be reduced for better performance of the solver.

5.1. Definition

Our ILP uses one integer variable $q_i \in \mathbb{Z}$ for every arc $a_i \in \mathcal{A}$ which represents the quantization of this arc. Every arc $a_i \in \mathcal{A}$ requires a non-negative quantization:

$$q_i \geq 0 \tag{1}$$

For consistency (cf. Section 4.1) we add the following constraint (analogous to previous work on consistent interval assignment for non-conforming partitions [CBK15, ULP*15]) for each pair of arc sets S and S_o which make up two opposite sides of a patch:

$$\sum_{a_i \in S} q_i - \sum_{a_j \in S_o} q_j = 0 \tag{2}$$

To ensure validity we add the following validity constraints which ensure that for each trace t_i one of its arcs between its origin and the first intersection n_{i*} with a trace starting in the $\pi/2$-sector around t_i is quantized to at least 1 (cf. Section 4.2):

$$\sum_{a_k \in S_{i*}} q_k \geq 1 \quad \forall \text{ traces } t_i \tag{3}$$

The layout constraints of Section 4.3 have a very similar form but are created for every intersection of two arcs forming a triangle with angles larger than α to prevent the creation of separatrices with excessive deviation:

$$\sum_{a_k \in S_{ji}} q_k \geq 1 \quad \forall\, n_{ij} \text{ with } \frac{l_{ji}}{l_{ij}} > \tan \alpha \tag{4}$$

Finally, to promote layout coarseness, we define the objective to be minimized as

$$E = \sum_{a_i \in \mathcal{A}} l_i^\perp \cdot q_i \;\; \rightarrow \min \tag{5}$$

where $l_i^\perp$ is half the parametric distance between the two arcs opposite of arc a_i (or half the parametric distance between the one opposite arc and a_i itself if a_i is boundary). Since the quantization of an arc specifies the number of quad strips that pass through this arc orthogonally (in the quad mesh implied by the quantized parametrization), this energy corresponds to the total length of the layout's quad strips.

An integer linear program is feasible if an assignment of variables exists such that all constraints are satisfied. Campen et al. [CBK15] show that a consistent quantization always exists in which all arcs are quantized to at least 1. Such a baseline quantization trivially fulfills all our constraints.

5.2. ILP Size

The size of the integer linear program described in the previous section depends largely on the size of the constructed T-mesh. For every arc there is one integer variable representing its quantized length (1), for every patch there are two consistency constraints (2), every trace adds a validity constraint (3), and for every node created


Figure 4: *On the* BUNNY *mesh with 192 traces the number of arcs quickly grows for decreasing α values to 189k at $\alpha = 0.25°$ but the number of integer variables remains around 252 (left). Since motorcycles need to be traced further and the resulting T-mesh consists of more elements for decreasing α, timings go up for both T-mesh construction and setup of the ILP (to 2.9 s and 0.4 s, respectively, at $\alpha = 0.25°$) but the time to solve the problem remains around 0.03 s.*

at an intersection a layout constraint (4) may be created. With decreasing angular bound α the traces get longer and the number of arcs, patches, and nodes in the T-mesh increases quickly (cf. Figure 4 left). In the following we describe how the program can be simplified significantly; interestingly, its size ultimately is proportional to the number of traces n, which is a constant independent of parameter α.

First of all, note that the number of T-junctions in the T-mesh is at most n since every trace can create at most one T-junction when it ends.

Now, we begin by looking at the number of variables and consistency constraints. Consider a simple strip of consecutive patches bounded by two traces, one on each side. Such a strip ends at T-junctions leading to wider patches (see inset top) or narrower patches (bottom). Running across the strip are individual arcs (blue) which need to be quantized to the same length for consistency (Equation (2)). By representing these arcs with the same variable, the consistency constraint is trivially fulfilled. Consistency constraints are then only needed for opposite patch sides in which at least one contains more than one arc which is only the case at T-junctions. Thus, the number of required constraints is at most n. Further, while strips tend to get longer for smaller values of α (cf. Figure 2) their total number, and thus also the total number of integer variables, is bounded from above by $\frac{3}{2}n$ as every strip starts and ends with a T-junction and a T-junction can be incident to at most three strips.

Finally, we consider the number of validity and layout constraints. Section 4.3 describes how the layout constraints are generated for a given trace when it intersects another one. These constraints define a set of arcs for which at least one has to be quantized to at least one. The sets of arcs contain all arcs between the start singularity and the intersection node. Similarly, the validity constraints also define a set of arcs between the start of a trace and an intersection which needs to be quantized to at least one (cf. Section 4.2). Thus, set S_{i*} and sets $\{S_{ij}\}$ created for trace i can be sorted by size, such that $S_0 \subset S_1 \subset \ldots \subset S_k$ as illustrated in the inset. Obviously, if an arc in S_0 is quantized to one, all supersets $S_1 \ldots S_k$ are quantized to at least one as well. Thus, a layout constraint needs to be created only from the smallest set per trace. Therefore, the numbers of integer variables, consistency constraints, validity constraints, and layout constraints are all $O(n)$.

6. Quad Mesh Extraction

A valid and consistent quantization such as produced by our algorithm specifies which singularities should lie on the same parametric iso-lines and thus defines which singularities should be connected by separatrices. There are a couple of ways to make use of this layout specification. For instance, one can generate a quad mesh adhering to this layout – in various ways.

In the context of the algorithm of [CBK15] our quantization could directly be used, as a drop-in replacement, to compute an integer grid map in which singularities are constrained onto specific integer locations resulting in a coarse quad mesh with one quad per layout patch. Alternatively, a path search on the T-mesh can be used to explicitly locate for each trace the closest singularity not separated in both coordinates. A separatrix between these two singularities could then be traced in the seamless input parametrization as in [RRP15, §6.1], obtaining an explicit embedding of the layout arcs.

For reasons of flexibility and guaranteed reliability, we opted to employ the recent re-parametrization approach of [LCBK19, §6] for our experiments. In a first step, this algorithm integrates the T-mesh into the underlying triangle mesh. Then, T-mesh re-embedding operations are used to get rid of all T-mesh arcs which are quantized to zero through collapsing. We then follow up with an additional step: we iteratively extend all T-junctions to the opposite


Figure 5: *Results of our algorithm on a variety of meshes. On the lower right model (*BRAIN*) a huge number (3.7K) of singularitites (i.e. layout nodes) are prescribed as input; while this naturally limits the level of coarseness that can be achieved, it serves to demonstrate the reliability and efficiency of our approach.*

Figure 6: *Higher values of α lead to layouts with a smaller number of patches #P.*

Model	#Faces	#Sing	α	#Traces	#Arcs	#Vars	#P	d_{mean}	d_{max}	MSJ$_{avg}$	t_{MCG}	t_{ILP}
ROCKERARM	20088	36	15°	144	2742	192	159	3.7°	14.7°	0.989	92 ms	30 ms
HEPTOROID	20000	166	10°	832	22822	1131	2051	2.3°	9.9°	0.982	348 ms	158 ms
SCULPT	7342	14	25°	88	936	107	54	5.5°	17.7°	0.982	29 ms	26 ms
KITTEN	100000	68	35°	272	2072	364	322	8.1°	34.9°	0.971	152 ms	43 ms
BOTIJO	29994	72	45°	320	2116	400	121	7.2°	41.4°	0.979	87 ms	34 ms
MASTERCYLINDER	100000	44	15°	192	3476	256	214	3.8°	14.2°	0.988	245 ms	42 ms
BLOCK	68352	48	35°	208	1604	291	76	6.6°	32.5°	0.985	116 ms	32 ms
BRAIN	100000	3721	45°	15332	122570	18722	23817	10.7°	45.0°	0.912	2214 ms	10522 ms
DUCK	19720	28	5°	104	5232	148	516	1.4°	3.8°	0.983	112 ms	37 ms
DUCK	19720	28	35°	104	852	138	77	9.8°	32.4°	0.976	38 ms	28 ms
TEST1	16323	10	15°	38	106	17	28	0.4°	6.0°	0.996	26 ms	22 ms
SPRAYER	21381	4	5°	22	93	10	12	0.1°	0.8°	0.998	26 ms	2 ms
GLUEGUN	12186	50	25°	244	1293	181	209	3.2°	24.2°	0.969	57 ms	31 ms
COGNIT	18934	54	25°	301	1631	199	181	2.1°	20.7°	0.988	66 ms	25 ms
CHAIN	5021	60	35°	303	1243	200	144	2.2°	34.6°	0.965	67 ms	27 ms
PUMP	2378	65	45°	320	1242	200	141	4.2°	40.7°	0.955	34 ms	28 ms
ENGINE	16502	24	15°	87	316	48	26	2.0°	10.1°	0.995	43 ms	25 ms
PART29	10698	12	35°	43	105	20	20	0.4°	3.1°	0.998	11 ms	3 ms

Table 1: *Statistical data for our results. From left to right: Model name, number of faces and singularities in the input, angular bound, number of traces, number of arcs in the motorcycle graph, number of variables in the reduced problem, number of patches in the resulting layout, mean and maximal separatrix deviation, average minimum scaled Jacobian of quads (of the depicted layout-aligned quad meshes), time for motorcycle graph construction and for solving of the ILP.*

sides of a patch, connecting opposing T-junctions if the quantization matches or splitting the corresponding opposite arc if not. The result is a T-mesh with no T-junctions left, explicitly representing all layout arcs (integrated into the triangle mesh as edge paths). Into each quad layout patch region we map a regular quad grid (size chosen compatibly based on the patches' parametric extent) by means of an optimized harmonic parametrization as described in [LCBK19, §6.2] for visualization purposes in the following.

7. Results

For the results in this section we created the input seamless parametrization by optimizing the energy proposed by Bommes et al. [BZK09] which minimizes the difference between the parametrization gradient to a cross field. The cross fields were obtained using the method of Bommes et al. [BZK09] with directional constraints as proposed by [CIE*16] (Sections 7 and 7.1), or using the method of Viertel et al. [VOS19] (Section 7.2). For models BOTIJO and ELK, as well as those in Section 7.3, we used the frame field provided in the supplemental material of [PPM*16]. We use Gurobi to solve the ILP and QEx [EBCK13] to extract the quad mesh from the final parametrization.

We show results of our method on a variety of models in Figure 5 and summarize statistical data in Table 1. Our layouts are typically well aligned and coarse – with α determining the balance. Both the construction of the T-mesh as well as solving the integer linear program typically take well below one second. Even on the BRAIN model with 3721 singularities the T-mesh construction completes in about 2 seconds and the ILP is solved in less than 11 seconds on a commodity PC, showing good scalability of our formulation.

Our algorithm is controlled by one parameter α which defines the maximum acceptable separatrix deviation from the seamless parametrization. As demonstrated in Figures 6 and 7 our layouts contain fewer patches at the cost of higher separatrix deviation for increasing values of α. Table 1 shows that the maximum separatrix deviation (d_{max}) generally stays below the given bound α, and the average deviation (d_{mean}) typically is significantly lower. Note that our formulation guarantees that the quad layout connectivity defined by the resulting quantization can always be embedded in the surface with deviation strictly bounded by α (e.g. by using the re-tracing strategy of [RRP15], cf. Section 6); when further optimizing or smoothing the embedding (e.g., following [LCBK19] or [CK14b]), this bound may be weakened of course.
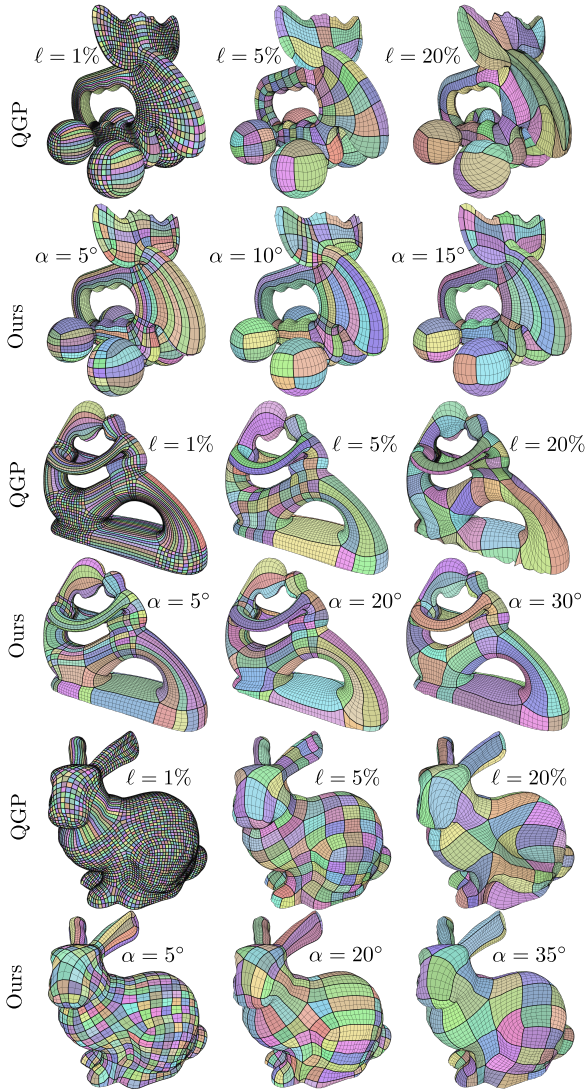
Figure 7: *Layout comparison. For QGP the target edge length ℓ is given in % of the bounding box diagonal.*

Model	Method	ℓ / α	#P	d_{mean}	d_{max}	MSJ_{avg}
ELK	CBK15	1%	3929	1.2°	33.3°	0.991
		5%	268	6.2°	71.6°	0.970
		20%	138	9.7°	82.9°	0.862
	ours	5°	412	1.1°	4.8°	0.983
		10°	206	2.2°	8.7°	0.986
		15°	134	3.2°	14.2°	0.982
FERTILITY	CBK15	1%	2109	1.3°	11.8°	0.994
		5%	161	5.9°	58.4°	0.980
		20%	103	9.2°	62.8°	0.880
	ours	5°	337	1.3°	4.3°	0.985
		20°	161	3.3°	17.7°	0.984
		30°	105	5.0°	29.3°	0.985
BUNNY	CBK15	1%	7033	0.8°	8.3°	0.991
		5%	333	7.5°	83.3°	0.971
		20%	146	10.1°	81.8°	0.891
	ours	5°	1088	1.5°	5.0°	0.977
		20°	344	4.3°	18.9°	0.971
		35°	148	7.5°	34.2°	0.958

Table 2: *Comparison with [CBK15]. In the third column we give the target edge length ℓ in % of the bounding box diagonal for [CBK15] and α for ours. #P, d_{mean}, d_{max}, and MSJ_{avg} are number of layout patches, mean and max separatrix deviation, and the average minimal scaled Jacobian of the quads, respectively.*


Figure 8: *Plotting number of layout patches #P against maximum (left) and mean (right) separatrix deviation. Our method (orange) consistently creates coarser layouts with less separatrix deviation than QGP (blue).*

7.1. Comparison with Campen et al.

We compare our results to those of the Quantized Global Parametrization (QGP) algorithm of Campen et al. [CBK15] who also create quad meshes via a quantization of a T-mesh. However, their quantization tries to minimize the difference between the arc *length* (rather than angle) in the input parametrization and the quantization which makes control over the resulting layout difficult. We generated quad meshes with their algorithm for target edge lengths ℓ of 1%, 5%, and 20% of the bounding box diagonal. The two coarser quad meshes are subdivided (via uniform scaling of the resulting parametrization) to achieve a similar number of quads as the result for 1% for better visual comparison. We applied our algorithm with α between 5° and 45° in steps of 5° and show results with the best matching number of layout patches. In Table 2 we report the number of layout patches, as well as average and maximum separatrix deviation.

For small target edge lengths leading to larger quantized values on the arcs QGP achieves results with uniformly sized quads but typically very dense layouts (Figure 7 left). With larger target edge lengths the base complex naturally becomes coarser as the quad mesh itself contains fewer quads (before subdivision). However, the separatrix deviation quickly increases and is often close to $\pi/2$, especially if the mesh contains regions with denser layout vertex distribution where the distance between pairs of layout vertices may be below the target edge length. Table 2 shows that QGP achieves lower element quality, as indicated by the average minimum scaled Jacobian, than our method except for the densest layout where the higher element quality comes at the cost of an order of magnitude more layout patches.

Low layout complexity and small separatrix deviation are competing goals. In Figure 8 we plot the number of layout patches against the maximum and mean separatrix deviation for different values of α and ℓ averaged over the models from Table 2. The plot as well as the numbers in Table 2 show that our algorithm consistently achieves a lower separatrix deviation with layouts of similar or smaller complexity.

Figure 9: *Our results on meshes with boundaries [VOS19].*

Figure 10: *Comparison with [PPM*16]. The color coding shows the angular deviation between input cross field and resulting parametrization, d is the area weighted average over the whole mesh.*

7.2. Comparison with Viertel et al.

Our algorithm shares some similarity with the one presented by Viertel et al. [VOS19] who create quad layouts (potentially non-conforming, containing T-junctions) by tracing a motorcycle graph and iteratively collapsing quad strips. In our algorithm assigning a quantization of 0 to an arc is equivalent to collapsing it [LCBK19]. Our ILP formulation has the advantage of finding the global optimum of a linear objective, Equation (5), and thus not suffering from running into local optima that a greedy optimization could reach and which Viertel et al. report may lead to T-junctions remaining in the layout which could have been avoided with a different collapse order. In addition, a valid and consistent quantization as produced by our algorithm always defines a conforming quad layout free of T-junctions. We present the result of our method on the models used in [VOS19] in Figure 9.

7.3. Comparision with Pietroni et al.

In Figure 10 we compare our results with those of [PPM*16]. Since Pietroni et al. create their layouts by searching separatrix candidates directly in a cross field, rather than a parametrization as in our case, we present the average deviation of the two directions induced by the resulting patchwise parametrization to the direction of the input cross field for a fair comparison. Note that, while we guarantee a maximal deviation to our (field-guided) input parametrization, the deviation compared to the guiding cross field may deviate more than α. The results show that, when aiming for a similar number of patches, we achieve lower distortion. For complex models Pietroni et al. generate layouts containing T-junctions. These can be converted to a conforming layout by extending all T-junctions. Doing this for the 78 T-junctions of the LION model, however, leads to a layout with over 82k patches. By contrast, our method yields a conforming layout with only 2.1k patches.

8. Conclusion & Future Work

We presented an algorithm for the creation of coarse quad layouts based on assigning integers to the arcs of a T-mesh by solving an

ILP. Our algorithm is controlled by only a single, intuitive parameter defining the maximum separatrix deviation from the seamless input parametrization. This could be especially useful in contexts where a certain minimum quality needs to be guaranteed while coarseness of the layout is also desirable. Our experiments show that the presented method not only achieves good results but is also computationally efficient, requiring only a couple of seconds for meshes with thousands of singularities.

Our algorithm relies on a seamless parametrization given as input to define layout singularities and desired separatrix alignment. As we base all our measurements on that seamless parametrization, our algorithm is oblivious to any metric distortion potentially already present in that parametrization. In cases where the parametrization is computed from a cross field – which may then be seen as the actual, original alignment intent – it is certainly tempting to omit the intermediate parametrization step. While a T-mesh similar to ours could directly be traced based on a cross field, as done in [MPZ14] and [PPM*16], a key challenge is dealing with matters of non-integrability and non-parametrizable singularity configurations in this setting. In any case, strict angle bounds will not be possible in a cross field based approach; arbitrarily large deviation may be necessary to obtain any valid layout in theory – unless T-junctions are acceptable. In that case rectangularity constraints could be implemented in a soft manner. The quantized T-mesh will then contain non-rectangular patches which could be resolved by inserting singularities [MPZ14] or by leaving T-junctions in the layout [PPM*16, VOS19].

In scenarios where strict adherence to the given input configuration is not required many interesting opportunities open up. For instance it would be interesting to explore variations of our approach that allow merging or splitting of singularities whenever beneficial for layout quality. It would also be interesting if the layout algorithm could make use of the fact that singularities may be free to move – at least within a certain range. In that case, small movements of layout vertices could move separatrices into the acceptable

bound, opening new opportunities to form a potentially coarser or better aligned layout.

Ultimately, for cases where the singularity configuration is entirely flexible, a holistic approach to layout creation would be desirable that, instead of separating the creation of layout vertices and the creation of their connecting separatrices in independent steps, creates both of them together in one step. This perspective is taken in [CK14a], albeit in a non-automatic approach, requiring user expertise and interaction.

Acknowledgements

References

[BCE*13] BOMMES D., CAMPEN M., EBKE H.-C., ALLIEZ P., KOBBELT L.: Integer-grid maps for reliable quad meshing. *ACM Transactions on Graphics 32*, 4 (2013), 98:1–98:12. 1, 3

[BLK11] BOMMES D., LEMPFER T., KOBBELT L.: Global structure optimization of quadrilateral meshes. *Computer Graphics Forum 30*, 2 (2011), 375–384. 2

[BLP*13] BOMMES D., LÉVY B., PIETRONI N., PUPPO E., SILVA C., TARINI M., ZORIN D.: Quad-mesh generation and processing: A survey. *Computer Graphics Forum 32*, 6 (2013), 51–76. 1

[BMRJ04] BOIER-MARTIN I. M., RUSHMEIER H. E., JIN J.: Parameterization of triangle meshes over quadrilateral domains. In *Proc. SGP '04* (2004), pp. 197–208. 2

[BZK09] BOMMES D., ZIMMER H., KOBBELT L.: Mixed-integer quadrangulation. *ACM Transactions on Graphics 28*, 3 (2009), 77:1–77:10. 1, 7

[Cam17] CAMPEN M.: Partitioning surfaces into quadrilateral patches: A survey. *Computer Graphics Forum 36*, 8 (2017), 567–588. 1

[CBK12] CAMPEN M., BOMMES D., KOBBELT L.: Dual Loops Meshing: Quality Quad Layouts on Manifolds. *ACM Transactions on Graphics 31*, 4 (2012), 110:1–110:11. 3

[CBK15] CAMPEN M., BOMMES D., KOBBELT L.: Quantized global parametrization. *ACM Transactions on Graphics 34*, 6 (2015). 3, 4, 5, 6, 8

[CIE*16] CAMPEN M., IBING M., EBKE H.-C., ZORIN D., KOBBELT L.: Scale-Invariant Directional Alignment of Surface Parametrizations. *Computer Graphics Forum 35*, 5 (2016). 7

[CK14a] CAMPEN M., KOBBELT L.: Dual strip weaving: Interactive design of quad layouts using elastica strips. *ACM Transactions on Graphics 33*, 6 (2014), 183:1–183:10. 10

[CK14b] CAMPEN M., KOBBELT L.: Quad layout embedding via aligned parameterization. *Computer Graphics Forum 33*, 8 (2014), 69–81. 7

[CSZZ19] CAMPEN M., SHEN H., ZHOU J., ZORIN D.: Seamless parametrization with arbitrary cones for arbitrary genus. *ACM Trans. Graph. 39*, 1 (2019). 2

[EBCK13] EBKE H.-C., BOMMES D., CAMPEN M., KOBBELT L.: QEx: Robust quad mesh extraction. 168:1–168:10. 7

[EGKT08] EPPSTEIN D., GOODRICH M. T., KIM E., TAMSTORF R.: Motorcycle Graphs: Canonical Quad Mesh Partitioning. *Computer Graphics Forum 27*, 5 (2008), 1477–1486. 3

[EH96] ECK M., HOPPE H.: Automatic reconstruction of B-spline surfaces of arbitrary topological type. In *Proc. SIGGRAPH 96* (1996), pp. 325–334. 2

[ESCK16] EBKE H.-C., SCHMIDT P., CAMPEN M., KOBBELT L.: Interactively controlled quad remeshing of high resolution 3d models. *ACM Trans. Graph. 35*, 6 (2016), 218:1–218:13. 1

[HSJ*20] HIEMSTRA R. R., SHEPHERD K. M., JOHNSON M. J., QUAN L., HUGHES T. J.: Towards untrimmed NURBS: CAD embedded reparameterization of trimmed B-rep geometry using frame-field guided global parameterization. *Computer Methods in Applied Mechanics and Engineering 369* (2020). 1

[KMZ11] KOVACS D., MYLES A., ZORIN D.: Anisotropic quadrangulation. *Computer Aided Geometric Design 28*, 8 (2011), 449–462. 1

[KNP07] KÄLBERER F., NIESER M., POLTHIER K.: Quadcover – surface parameterization using branched coverings. *Computer Graphics Forum 26*, 3 (2007). 1

[LCBK19] LYON M., CAMPEN M., BOMMES D., KOBBELT L.: Parametrization quantization with free boundaries for trimmed quad meshing. *ACM Trans. Graph. 38*, 4 (2019). 3, 6, 7, 9

[MAB*19] MARINOV M., AMAGLIANI M., BARBACK T., FLOWER J., BARLEY S., FURUTA S., CHARROT P., HENLEY I., SANTHANAM N., FINNIGAN G. T., MESHKAT S., HALLET J., SAPUN M., WOLSKI P.: Generative design conversion to editable and watertight boundary representation. *Computer-Aided Design 115* (2019), 194 – 205. 1

[MC19] MANDAD M., CAMPEN M.: Exact constraint satisfaction for truly seamless parametrization. *Computer Graphics Forum 38*, 2 (2019), 135–145. 3

[MK12] MÖBIUS J., KOBBELT L.: OpenFlipper: An open source geometry processing and rendering framework. In *Curves and Surfaces*, vol. 6920 of *Lecture Notes in Computer Science*. 2012. 10

[MPZ14] MYLES A., PIETRONI N., ZORIN D.: Robust field-aligned global parametrization. *ACM Transactions on Graphics 33*, 4 (2014). 9, 10

[PPM*16] PIETRONI N., PUPPO E., MARCIAS G., SCOPIGNO R., CIGNONI P.: Tracing field-coherent quad layouts. *Computer Graphics Forum 35*, 7 (2016). 2, 7, 9, 10

[PPTSH14] PANOZZO D., PUPPO E., TARINI M., SORKINE-HORNUNG O.: Frame fields: Anisotropic and non-orthogonal cross fields. *ACM Transactions on Graphics 33*, 4 (2014), 134. 1

[RP17] RAZAFINDRAZAKA F. H., POLTHIER K.: Optimal base complexes for quadrilateral meshes. *Computer Aided Geometric Design 100*, 52-53 (2017), 63–74. 2

[RRP15] RAZAFINDRAZAKA F. H., REITEBUCH U., POLTHIER K.: Perfect matching quad layouts for manifold meshes. *Computer Graphics Forum 34*, 5 (2015), 219–228. 2, 6, 7

[TPP*11] TARINI M., PUPPO E., PANOZZO D., PIETRONI N., CIGNONI P.: Simple quad domains for field aligned mesh parametrization. *ACM Transactions on Graphics 30*, 6 (2011). 1, 2

[ULP*15] USAI F., LIVESU M., PUPPO E., TARINI M., SCATENI R.: Extraction of the quad layout of a triangle mesh guided by its curve skeleton. *ACM Transactions on Graphics 35*, 1 (2015), 6. 5

[VCD*16] VAXMAN A., CAMPEN M., DIAMANTI O., PANOZZO D., BOMMES D., HILDEBRANDT K., BEN-CHEN M.: Directional field synthesis, design, and processing. *Computer Graphics Forum 35*, 2 (2016). 2

[VOS19] VIERTEL R., OSTING B., STATEN M.: Coarse quad layouts through robust simplification of cross field separatrix partitions. *Proc. 28th International Meshing Roundtable* (2019). 2, 7, 9, 10

[ZZY16] ZHANG S., ZHANG H., YONG J.-H.: Automatic quad patch layout extraction for quadrilateral meshes. *Computer-Aided Design and Applications 13*, 3 (2016), 409–416. 2