

Conformal Factor Persistence for Fast Hierarchical Cone Extraction

A. M. Vintescu¹, F. Dupont², G. Lavoue², P. Memari³ and J. Tierny⁴

¹LTCI - Télécom ParisTech, Paris, France

²LIRIS - Université de Lyon, CNRS, Lyon, France

³CNRS, LIX - Ecole Polytechnique, Palaiseau, France

⁴Sorbonne Universités, UPMC Univ Paris 06, CNRS, LIP6 UMR 7606, France

Abstract

This paper presents a new algorithm for the fast extraction of hierarchies of cone singularities for conformal surface parameterization. Cone singularities have been shown to greatly improve the distortion of such parameterizations since they locally absorb the area distortion. Therefore, existing automatic approaches aim at inserting cones where large area distortion can be predicted. However, such approaches are iterative, which results in slow computations, even often slower than the actual subsequent parameterization procedure. This becomes even more problematic as often the user does not know in advance the right number of needed cones and thus needs to explore cone hierarchies to obtain a satisfying result. Our algorithm relies on the key observation that the local extrema of the conformal factor already provide a good approximation of the cone singularities extracted with previous techniques, while needing only one linear solving where previous approaches needed one solving per hierarchy level. We apply concepts from persistent homology to organize very efficiently such local extrema into a global hierarchy. Experiments demonstrate the approximation quality of our approach quantitatively and report time-performance improvements of one order of magnitude, which makes our technique well suited for interactive contexts.

1. Introduction

Surface parameterization [FH05, SPR06] is a fundamental tool for numerous computer graphics and geometry processing applications, as it enables the unfolding of triangulated surfaces embedded in 3D to simpler geometrical domains, such as the 2D plane, where geometrical tasks can be carried out more efficiently or in a simpler manner. Popular applications of surface parameterization include texture mapping, morphing, detail transfer, mesh completion and re-meshing. To be unfolded to the plane, an input surface must have the topology of a disk or alternatively, must be artificially decomposed into charts with disk topology. While in the first case the unfolding will usually be accompanied with significant area distortion, in the second case, the unfolding will additionally exhibit visual artifacts due to parameterization discontinuities across the boundaries of the charts. To overcome these two issues, several techniques have been proposed for the *global parameterization* of triangulated surfaces [RLL*06, KSS06, BCGB08, SSP08, YGL*09, BZK09, MZ12, MZ13]. The global parameterization of a surface with disk topology can be defined as a homeomorphism from the surface to a subset of the plane, such that the discrete Gaussian curvature is zero everywhere except for a few vertices called *cone singularities*.

Several approaches based on metric scaling have been proposed in the past to address global parameterization [KSS06, BCGB08, SSP08, YGL*09, MZ12, MZ13]. These algorithms traditionally take as an additional input a set of cone singularities through which the surface will be cut to be given disk topology. Several algorithms [BCGB08, SSP08, MZ12, MZ13] have been proposed for their automatic extraction, as their number and location have a dras-

tic effect on the distortion of the output parameterization. These techniques aim at inserting cones where large area distortion can be predicted prior to the actual parameterization. In particular, they rely on the solving of a sequence of linear systems, to iteratively insert cones (one solving per hierarchy level). However, this results in slow computations, even often slower than the actual subsequent parameterization procedure. This becomes even more problematic as often the user does not know in advance the right number of needed cones and thus needs to explore cone hierarchies to obtain a satisfying result.

This paper addresses this problem by introducing a novel technique for the fast computation and exploration of hierarchies of cone singularities for conformal surface parameterization. Our approach relies on the key observation that the local extrema of the conformal factor already provide a good approximation of the cone singularities extracted with previous techniques, while needing only one linear solving where previous approaches needed one solving per hierarchy level. We apply concepts from persistent homology [ELZ00, EH09] to organize very efficiently such local extrema into a global hierarchy. Experiments demonstrate the approximation quality of our approach quantitatively. In particular, the analysis of the area distortion of the parameterizations obtained by using the cones extracted with our approach shows comparable (if not improved) distortion when compared to the cones extracted by an iterative technique [BCGB08], while yielding time-performance improvements of one order of magnitude, which makes our technique well suited for interactive contexts.

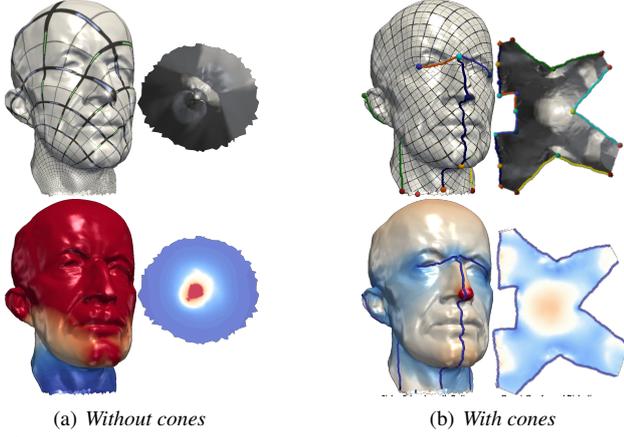


Figure 1: (a) The parameterization of a surface using a standard conformal approach [LPRM02] results in high average area distortion ($D_a = 47.361$). (b) Cone singularities (here manually inserted), used in conjunction with a metric-scaling parameterization approach, locally absorb such a distortion and drastically reduce it globally ($D_a = 1.487$). Top to bottom: (i) textured surface with cones (spheres) and conic cuts (cylinders), as well as planar unfolding, (ii) area distortion on the surface and in the planar domain (color map: blue (0) to red (10), ideal value: 1).

2. Preliminaries

The input surface M is given as a mesh made of vertices (V), edges (E) and triangles (T). Their number is noted with $|V|$, $|E|$ and $|T|$ respectively. The geometry of M is given as the 3D coordinates of the vertices $X_v = (v_x, v_y, v_z), \forall v \in V$. The output parameterization is represented with 2D coordinates for each vertex $U_v = (u, v)$. The length of an edge is given by $e_{i,j} = \|X_{v_i} - X_{v_j}\|_2$ in 3D or $e_{i,j} = \|U_{v_i} - U_{v_j}\|_2$ in 2D. An angle in a triangle t is given by: $\alpha_{v_i}^t = \arccos\left(\frac{e_{i,j}^2 + e_{i,k}^2 - e_{j,k}^2}{2e_{i,j}e_{i,k}}\right)$, where v_i, v_j and v_k are the vertices of t . The discrete Gaussian curvature is given by $K = \left\{ k_{v_i} = \begin{cases} 2\pi - \sum_{t \in T_{v_i}} (\alpha_{v_i}^t), & \text{for an interior vertex} \\ \pi - \sum_{t \in T_{v_i}} (\alpha_{v_i}^t), & \text{for a boundary vertex} \end{cases} \right\}$, where T_{v_i} represents the set of incident triangles to the vertex v_i . The Gauss-Bonnet Theorem states that the integral of the curvature is a constant, which depends on the topology of M : $\sum K = 2\pi\chi$, where χ represents the Euler characteristic of M ($\chi = |V| - |E| + |T|$). Given a surface with disk topology, a *global* parameterization can be defined as a conformal (i.e. angle preserving) homeomorphism to a subset of the plane, such that the discrete Gaussian curvature is zero everywhere except at a set of selected vertices C , called *cones*.

To be unfolded to the plane, the input surface must have disk topology. To this effect surfaces with sphere topology are cut along *conic cuts* – shortest paths that connect the cones (Fig. 1(b)). We detail this process hereafter. Variants of this strategy can be derived for surfaces with different genus.

The shortest path between each possible pair of cones is first computed with Dijkstra’s algorithm. Next, a minimum spanning tree is constructed on a graph where the nodes denote the cones and where each edge is weighted by the geodesic distance between its cones (i.e. the length of their shortest paths). Given this spanning tree, the surface is then cut open along the shortest paths cor-

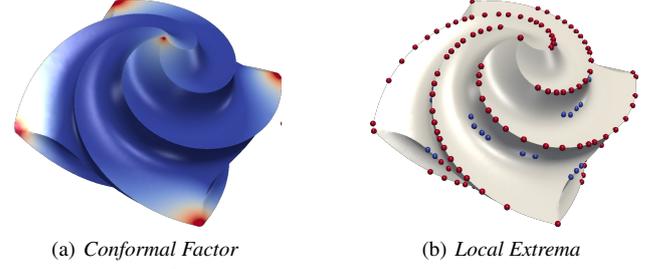


Figure 2: (a) Conformal factor of a surface (from blue to red, normalized to $[0, 1]$). (b) Local minima (blue spheres), respectively maxima (red spheres), are extracted as vertices such that all neighbors have a higher, respectively lower, conformal factor. The key observation of our approach is that the cone singularities extracted with previous approaches often constitute a subset of these extrema.

responding to its edges, and we call such paths *conic cuts*. Once the mesh is cut open along the conic cuts, a metric scaling approach [BCGB08, SSP08, YGL*09, MZ12, MZ13] can be used to compute a conformal parameterization that has zero discrete Gaussian curvature everywhere except at the cones. The benefit of such a strategy is shown in Fig. 1, which compares the area distortion of a standard conformal parameterization (Fig. 1(a)) with that of an approach based on cone singularities and metric scaling (Fig. 1(b)).

3. Conformal Factor Persistence

This section presents our new algorithm for the fast computation of hierarchies of cone singularities for conformal parameterizations with low area distortion.

3.1. Conformal Factor Computation

As shown in Fig. 1, cones can be interpreted as *distortion absorbers*. Thus, our strategy consists in placing cones where large distortion can be predicted. We achieve this by evaluating the *conformal factor*. While previous approaches relied on an iterative introduction of cones at the global minimum and maximum of conformal factors for varying target curvatures [BCGB08, SSP08, MZ12], we observed that in practice, most of the cones identified with these approaches are actually already present as the most persistent local extrema of the conformal factor at the *first* iteration. Our approach relies on this key observation.

Given a triangular surface, two metrics (i.e. edge lengths) are said to be discretely conformally equivalent if they are related through a *discrete scaling factor* [SSP08, BCGB08, BCG08]. This factor can be found as a solution to the Poisson equation $\nabla^2 \phi = K^T - K^O$, where ∇^2 is the cotangent weights Laplacian and where K^O and K^T are the original and target curvatures respectively. Deriving from the Gauss-Bonnet theorem, $\sum_{v_i \in V} (k_{v_i}^O) = \sum_{v_i \in V} (k_{v_i}^T)$. The conformal factor is obtained after solving the Poisson equation by assigning to each vertex a target curvature that depends on the normalized area of its incident triangles:

$$k_{v_i}^T = \sum_{v_i \in V} (k_{v_i}^O) \left(\frac{\sum_{f \in T_{v_i}} (A_f/3)}{A_M} \right)$$

where A_f represents the area of triangle f and $A_M = \sum_{f \in T} (A_f)$ represents the total area of the surface. The resulting conformal factor indicates the amount of “work” that is required to unfold a surface onto a domain with constant curvature (see Fig. 2).

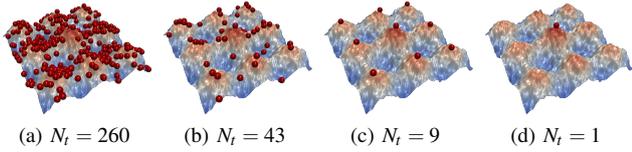


Figure 3: Hierarchy of persistent maxima (red spheres) on a noisy terrain (elevation map from blue to red), composed of 8 hills surrounding a higher, central hill. The hierarchy progressively discards small bumps until the major peaks are identified (c). Further, the hierarchy sorts the major peaks in function of their importance and retrieves the highest peak at the coarsest hierarchy level (d).

3.2. Persistence-driven Extremum Selection

As shown in Fig. 2, the local extrema of the conformal factor ϕ computed previously are located in configurations of large, local variations in conformal factor. The key observation of our approach is that the cone singularities extracted with previous approaches often constitute a subset of these extrema. We employ concepts from persistent homology [ELZ00, EH09] to efficiently organize these extrema into a cone hierarchy that can be interactively explored.

Persistent critical point pairs can be easily extracted by tracking the connectivity of the sub- and sur-level sets of the conformal factor. The vertices are first visited in order of increasing conformal factor and connected components of visited vertices are tracked by a Union-Find (UF) data-structure [CLRS01]. If a visited vertex v has no other visited vertices among its neighbors, this means v is a local minimum of ϕ : a new set is created in the UF data-structure and it is associated to v . Otherwise, v is associated to the UF set of its visited neighbors. If multiple such sets exist, this means that v is a saddle of ϕ that joins distinct connected components of sub-level sets. In such a case, v is paired with the minimum m associated to the *youngest* of these UF sets (to the one that has been created last). This pairing strategy is known as the Elder’s rule [EH09]. The absolute difference $p(m, v) = |\phi(m) - \phi(v)|$ is called the *persistence* of the pair (m, v) . Once the entire set of vertices has been visited, each local minimum has been paired exactly once and has therefore a *persistence* value. The same procedure is applied symmetrically with the maxima, by visiting the vertices in order of decreasing ϕ . The persistence of each extremum naturally defines a hierarchy. In particular, to explore such a cone hierarchy, the user provides a target number N_t and we alternatively add a maximum and a minimum to the hierarchy, in decreasing order of persistence.

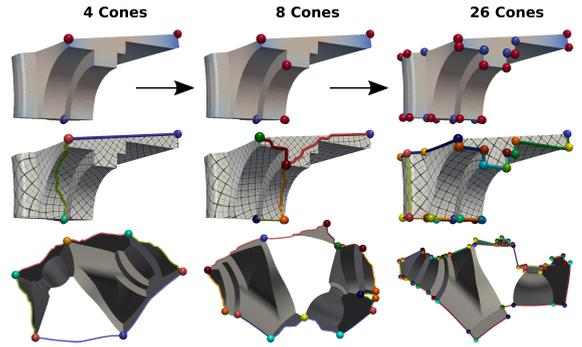
As shown in Fig. 3, persistence driven hierarchies tend to organize extrema in function of their importance. As a result, the most persistent extrema of the conformal factor will be located at the extremities of the most prominent features of the shape (Fig. 4).

The computation of the persistence pairs is very efficient in practice. It only requires $O(n \log(n) + N\alpha(N))$ steps, where n and N are the number of vertices and simplices in the mesh respectively and where $\alpha(\cdot)$ is the inverse of the Ackermann function (thus, $\alpha(\cdot)$ is an extremely slow-growing function).

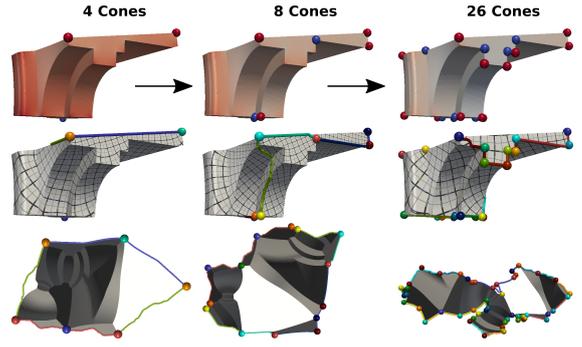
4. Results

This section reports experiments performed with a C++ implementation of our approach on a laptop with a 2.5 GHz i7 CPU.

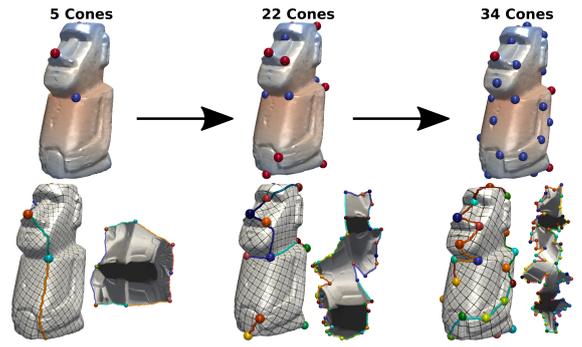
Fig. 4 presents a visual comparison between the cone hierarchies



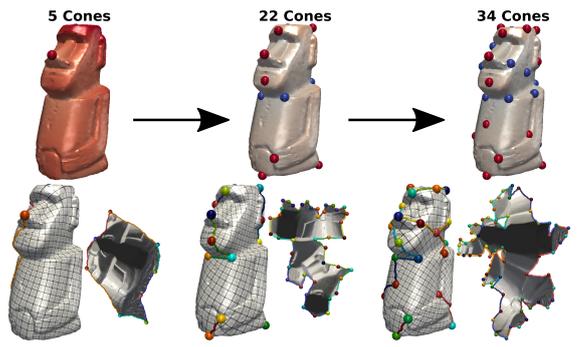
(a) Fandisk Model - Our Approach



(b) Fandisk Model - Iterative Approach [BCGB08]



(c) Moai Model - Our Approach



(d) Moai Model - Iterative Approach [BCGB08]

Figure 4: Visual comparison between the cone hierarchies extracted with our approach and an iterative approach [BCGB08]. Top: Conformal factor (blue to red) and cone singularities, middle and bottom: surface parameterization and unfolding after metric-scaling. Note that the cones extracted with our approach are visually similar to those extracted by an iterative approach.

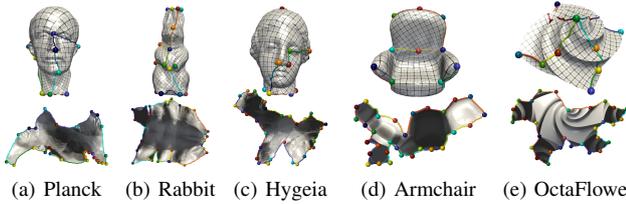


Figure 5: Conformal parameterizations based on metric scaling using the cones extracted with our approach (top: surface parameterization, bottom: unfolded surface).

extracted with our algorithm and the iterative approach by Ben-Chen et al. [BCGB08], as it is probably the iterative method which is the most similar in spirit to ours. That approach iteratively inserts pairs of cones as the global minimum and maximum of conformal factor evaluations (one evaluation per hierarchy level). In particular, notice on the first row that the cones extracted with our approach are visually similar to those extracted with the iterative approach. Tab. 1 provides a quantitative comparison with our approach to cone detection and an iterative one, for the examples shown in Fig. 4 and 5 (which constitutes a reasonably diversified collection of organic and mechanical shapes). To quantitatively inspect the quality of our approach, we compare the output areal and quasi-conformal distortion yielded by the same conformal parameterization procedure based on metric scaling, applied on the two sets of cones (last 4 columns). These evaluations show that the quasi-conformal distortion is nearly identical while the areal distortion is improved with our technique in all but two examples. Since our approach relies on only one evaluation of the conformal factor, it requires only one linear solving for the computation of the entire hierarchy. Thus, it is much faster in practice than an iterative approach [BCGB08] (23 times on average). Note that the timings reported for the iterative approach correspond to the time required to complete the queried hierarchy level, while the timings reported for our approach correspond to the entire hierarchy extraction. Finally, note that, in contrast to an iterative approach, our algorithm always takes less time to extract the entire hierarchy than to actually parameterize the surface, which makes it well suited for an interactive exploration of the cone configurations.

5. Conclusion

This paper presented a new algorithm for the fast extraction of cone hierarchies for conformal surface parameterization. Our approach is based on the key observation that the local extrema of the con-

Table 1: Quantitative comparison between metric scaling based parameterizations obtained with cones extracted with our approach and an iterative one [BCGB08]. On average, our cone detection strategy is 23 times faster than an iterative approach, while yielding comparable if not improved output distortions. $|V|$ and $|C|$ stand for the number of vertices and cones respectively.

Model	$ V $	$ C $	Time (s)				Distortion			
			Cone Detection		Parameterization		Areal		Quasi Conformal	
			Ours	[BCGB08]	Ours	[BCGB08]	Ours	[BCGB08]	Ours	[BCGB08]
Planck (Fig. 5(a))	28K	16	0.246	5.842	2.793	1.550	1.559	1.199	1.018	1.013
Armchair (Fig. 5(d))	40K	16	0.265	6.147	2.385	2.754	1.396	1.513	1.014	1.015
Rabbit (Fig. 5(b))	57K	16	0.312	10.077	4.184	4.277	1.113	1.172	1.009	1.009
OctaFlower (Fig. 5(e))	32K	16	0.184	5.896	1.638	2.396	1.017	1.117	1.009	1.008
Hygeia (Fig. 5(c))	33K	16	0.188	4.404	1.441	1.174	1.125	1.184	1.021	1.022
Fandisk (Fig. 4)	29K	4	0.211	1.454	1.997	1.888	1.578	1.611	1.017	1.015
Fandisk (Fig. 4)	29K	8	-	2.767	2.064	1.910	1.426	1.583	1.013	1.012
Fandisk (Fig. 4)	29K	26	-	6.764	1.888	1.874	1.371	1.587	1.008	1.012
Moai (Fig. 4)	10K	5	0.055	0.329	0.281	0.249	1.285	1.233	1.043	1.041
Moai (Fig. 4)	10K	22	-	1.023	0.248	0.247	1.227	1.231	1.034	1.033
Moai (Fig. 4)	10K	34	-	1.497	0.238	0.249	1.171	1.249	1.036	1.034

formal factor represent good cone candidates and that persistent homology concepts can help for their efficient organization into a hierarchy. Although computationally efficient, our approach can be seen as an approximation of iterative methods, as it may not deliver the exact same output (Fig. 4). However, experiments still demonstrated quantitatively the relevance of our approach in terms of output distortions and computation speedups. Also, despite their geometrical relevance, the cones provided by our approach may not coincide perfectly with locations desired by an artist in interactive parameterization tasks. We believe they should rather be considered as fast initial guesses in that context. In the future, we would like to investigate how our cone hierarchies could be used as surface feature points [LF09, BBB*11] for surface correspondence and cross parameterization purposes.

Acknowledgments

This work is supported by the ANR (Agence Nationale de la Recherche) through the CrABEx project (ANR-13-CORD-0013).

References

- [BBB*11] BOYER E., BRONSTEIN A., BRONSTEIN M., BUSTOS B., DAROM T., HORAUD R., HOTZ I., KELLER Y., KEUSTERMANS J., KOVNATSKY A., LITMAN R., REININGHAUS J., SPIRAN I., SMEETS D., SUETENS P., VANDERMEULEN D., ZAHARESCU A., ZOBEL V.: SHREC 2011: Robust feature detection and description benchmark. In *Eurographics Workshop on 3D Object Retrieval (3DOR)* (2011). 4
- [BCG08] BEN-CHEN M., GOTSMAN C.: Characterizing Shape Using Conformal Factors. In *Eurographics 2008 Workshop on 3D Object Retrieval* (2008). 2
- [BCGB08] BEN-CHEN M., GOTSMAN C., BUNIN G.: Conformal flattening by curvature prescription and metric scaling. *Comp. Graph. Forum* 27, 2 (2008), 449–458. 1, 2, 3, 4
- [BZK09] BOMMES D., ZIMMER H., KOBELT L.: Mixed-integer quadrangulation. *ACM Trans. Graph.* (2009). 1
- [CLRS01] CORMEN T., LEISERSON C., RIVEST R., STEIN C.: *Introduction to Algorithms*, 2nd ed. MIT Press, (2001). 3
- [EH09] EDELSBRUNNER H., HARER J.: *Computational Topology: An Introduction*. American Mathematical Society, (2009). 1, 3
- [ELZ00] EDELSBRUNNER H., LETSCHER D., ZOMORODIAN A.: Topological Persistence and Simplification. In *Proc. of FOCS* (2000). 1, 3
- [FH05] FLOATER M. S., HORMANN K.: Surface parameterization: a tutorial and survey. In *Adv. Multires. Geom. Mod.* Springer, (2005). 1
- [KSS06] KHAREVYCH L., SPRINGBORN B., SCHRÖDER P.: Discrete conformal mappings via circle patterns. *ACM Trans. Graph.* (2006). 1
- [LF09] LIPMAN Y., FUNKHOUSER T.: Möbius voting for surface correspondence. *ACM Trans. Graph.* (2009). 4
- [LPRM02] LÉVY B., PETITJEAN S., RAY N., MAILLOT J.: Least squares conformal maps for automatic texture atlas generation. *ACM Trans. Graph.* 21, 3 (2002), 362–371. 2
- [MZ12] MYLES A., ZORIN D.: Global parametrization by incremental flattening. *ACM Trans. Graph.* 31, 4 (2012), 109:1–109:11. 1, 2
- [MZ13] MYLES A., ZORIN D.: Controlled-distortion constrained global parametrization. *ACM Trans. Graph.* 32, 4 (2013), 1–14. 1, 2
- [RLL*06] RAY N., LI W. C., LÉVY B., SHEFFER A., ALLIEZ P.: Periodic global parameterization. *ACM Trans. Graph.* 25, 4 (2006). 1
- [SPR06] SHEFFER A., PRAUN E., ROSE K.: Mesh parameterization methods and their applications. *Found. Trends. Comp. Graph. Vis.* 2, 2 (2006), 105–171. 1
- [SSP08] SPRINGBORN B., SCHRÖDER P., PINKALL U.: Conformal equivalence of triangle meshes. *ACM Trans. Graph.* (2008). 1, 2
- [YGL*09] YANG Y. L., GUO R., LUO F., HU S. M., GU X.: Generalized discrete ricci flow. *Comp. Graph. Forum* (2009). 1, 2