Paneling Architectural Freeform Surfaces

Michael Eigensatz
ETH Zurich / EPFL

Martin Kilian Alexander Schiftner
Evolute / TU Wien

Niloy J. Mitra
KAUST / IIT Delhi

Helmut Pottmann
KAUST / TU Wien

Mark Pauly
EPFL

Figure 1: Rationalization of large-scale architectural freeform surfaces with planar, single-, and double-curved panels. Our algorithm computes a paneling solution that meets prescribed thresholds on positional and normal continuity, while minimizing total production cost. Reuse of molds (left) and predominant use of simple panels (right) are important drivers of the optimization. (Left: Zaha Hadid Architects, Lilium Tower, Warsaw. Right: Zaha Hadid Architects, National Holding Headquarters, Abu Dhabi.)

Abstract

The emergence of large-scale freeform shapes in architecture poses big challenges to the fabrication of such structures. A key problem is the approximation of the design surface by a union of patches, so-called panels, that can be manufactured with a selected technology at reasonable cost, while meeting the design intent and achieving the desired aesthetic quality of panel layout and surface smoothness. The production of curved panels is mostly based on molds. Since the cost of mold fabrication often dominates the panel cost, there is strong incentive to use the same mold for multiple panels. We cast the major practical requirements for architectural surface paneling, including mold reuse, into a global optimization framework that interleaves discrete and continuous optimization steps to minimize production cost while meeting user-specified quality constraints. The search space for optimization is mainly generated through controlled deviation from the design surface and tolerances on positional and normal continuity between neighboring panels. A novel 6-dimensional metric space allows us to quickly compute approximate inter-panel distances, which dramatically improves the performance of the optimization and enables the handling of complex arrangements with thousands of panels. The practical relevance of our system is demonstrated by paneling solutions for real, cutting-edge architectural freeform design projects.

Keywords: architectural geometry, freeform design, rationalization, geometric optimization

ACM Reference Format

Eigensatz, M., Kilian, M., Schiftner, A., Mitra, N., Pottmann, H., Pauly, M. 2010. Paneling Architectural Freeform Surfaces. *ACM Trans. Graph.* 29, 4, Article 45 (July 2010), 10 pages. DOI = 10.1145/1778765.1778782 <http://doi.acm.org/10.1145/1778765.1778782>

Copyright Notice

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, fax +1 (212) 869-0481, or permissions@acm.org.
© 2010 ACM 0730-0301/2010/07-ART45 \$10.00 DOI 10.1145/1778765.1778782
<http://doi.acm.org/10.1145/1778765.1778782>

1 Introduction

Freeform shapes play an increasingly important role in contemporary architecture. With the emergence of large-scale architectural freeform surfaces the essential question arises of how to proceed from a geometrically complex design towards a feasible and affordable way of production. This fundamental problem, in the architectural community referred to as *rationalization*, is largely related to the issue of *paneling*, i.e., the segmentation of a shape into simpler surface patches, so-called panels, that can be fabricated at reasonable cost with a selected manufacturing process (see Figure 1). The paneling problem can arise both for the exterior and interior skin of a building, and plays a central role in the design specification phase of any architectural project involving freeform geometry.

Recent technological advances enable the production of single- and double-curved panels that allow a faithful approximation of curved surfaces. While planar panels are always the most cost-effective, the progression towards the more expensive general freeform panels depends on the panel material and manufacturing process. Most commonly, curved panels are produced using molds with the cost of mold fabrication often dominating the panel cost (see Figure 9). There is thus a strong incentive to reuse the same mold for the production of multiple panels to reduce the overall cost.

Our goal is to find a paneling solution for a given freeform design that achieves prescribed quality requirements, while minimizing production cost and respecting application-specific constraints. The *quality* of the paneling is mainly determined by the geometric closeness to the input surface, the positional and normal continuity between neighboring panels, and the fairness of corresponding panel boundary curves. The *cost* mostly depends on the size and number of panels, the complexity of the panel geometry, and the degree of reuse of molds that need to be custom-built to fabricate the panels. A key objective of our work is to solve instances of the paneling problem on large-scale architectural freeform designs that often consist of thousands of panels. Due to the high complexity and global coupling of optimization objectives and constraints, manual layout of panels for these freeform surfaces is infeasible, mandating the use of advanced computational tools.

1.1 Contributions

We introduce a computational approach to freeform surface paneling. As our main contributions, we

- identify the key aspects of the paneling problem that are amenable to computation and derive a mathematical framework that captures the essential design goals,
- present an algorithmic solution based on a novel global optimization method that alternates between discrete and continuous optimization steps to improve the quality of the paneling while reducing cost through mold reuse,
- introduce a novel 6-dimensional metric space to allow fast computation of approximate inter-panel distances, which dramatically improves the performance of the optimization and enables the handling of complex arrangements with thousands of panels, and
- demonstrate the practical relevance of our system by computing paneling solutions for real, cutting-edge designs that currently cannot be realized at the desired aesthetic quality.

1.2 Related Work

Early contributions to the field of freeform architecture come from research at Gehry Technologies (see e.g. [Shelden 2002]). These are mostly dedicated to developable or nearly developable surfaces, as a result of the specific design process that is based on digital reconstruction of models made from material that assumes (nearly) developable shapes.

Research on freeform architecture is promoted by the Smart Geometry group (www.smartgeometry.com), whose interest so far mostly focussed on parametric design tools. These can be helpful for shape generation processes that have panel properties built into them. However, such a forward approach makes it very difficult to achieve the desired shapes and obtain a satisfactory paneling solution for sufficiently complex geometries.

Most previous work on the paneling problem deals with planar panels. Initial research in this direction dealt with special surface classes [Glymph et al. 2002]. Covering general freeform surfaces with planar quad panels could be approached with methods of discrete differential geometry [Bobenko and Suris 2009] and led to new ways of supporting beam layout and the related computation of multi-layer structures [Liu et al. 2006; Pottmann et al. 2007]. More recently, this approach was extended to the covering of freeform surfaces by single-curved panels arranged along surfaces strips [Pottmann et al. 2008b]. Additional results in this direction, e.g., hexagonal meshes with planar faces, have been presented at “Advances in Architectural Geometry” [Pottmann et al. 2008a]. The idea of optimizing for repeated elements by altering the vertex positions of a given mesh is explored by Fu et al. [2010] in the context of quad meshes and by Singh and Schaefer [2010] in the context of triangle meshes, in order to create a set of reusable pre-fabricated tiles.

Our approach bears some similarity to variational methods for approximating a surface with simple geometric primitives. Originally introduced by Cohen-Steiner et al. [2004] for surface approximation by planes, various extensions have been proposed for additional surface types, e.g., spheres and cylinders [Wu and Kobbelt 2005], quadrics [Yan et al. 2006], or developable surfaces [Julius et al. 2005]. Recently, an optimization has been proposed to simultaneously partition the input surface, as well as determine the types and number of shape proxies required [Li et al. 2009]. These

Figure 2: Projects involving double-curved panels where a separate mold has been built for each panel. These examples illustrate the importance of the curve network and the challenges in producing architectural freeform structures. (Left: Peter Cook and Colin Fournier, Kunsthau, Graz. Right: Zaha Hadid Architects, Hungerburgbahn, Innsbruck.)

methods optimize for a surface segmentation to reduce the approximation error. In our setting, the segmentation is part of the design specification and we optimize for position and tangent continuity across panel boundaries, allowing systematic deviations from the reference surface to improve the paneling quality and reduce cost. Enabling mold reuse and aesthetic control, which are key requirements of architectural rationalization, necessitates a substantially different approach both in the underlying formulation of the optimization as well as its implementation. Similarly, state-of-the-art methods in surface fitting and local registration (see e.g. [Varady and Martin 2002; Shamir 2008]), while an integral component of our system, are insufficient to solve large-scale freeform paneling problems.

In shape analysis, the problems of symmetry detection [Mitra et al. 2006; Podolak et al. 2006] and regularity detection [Pauly et al. 2008] involve identification and extraction of repeated elements, exact and approximate, in 3D geometry. Subsequently, detected repetitions can be made exact by symmetrization using subtle modifications of the underlying meshing structure [Mitra et al. 2007; Golovinskiy et al. 2009] thus deforming a surface towards an exact symmetric configuration. These methods, designed to enhance detected symmetries, are unsuited for handling architectural freeform designs where large repeated sections are exceptions rather than the norm. Our optimization has a similar symmetrizing effect in enabling trading approximation error for a stronger degree of mold reuse, leading to significant savings in terms of manufacturing cost.

2 Problem Specification

In this section we introduce a specification of the paneling problem and the corresponding terminology common in architectural design (see Figure 3). The specification is the result of extensive consultations with architects and our experience with real-world freeform design projects, some of which we highlight in Section 4.

2.1 Terminology

Panels and molds. Let F be the given input freeform surface describing the shape of the design. Our goal is to find a collection $\mathcal{P} = \{P_1, \dots, P_n\}$ of panels P_i , such that their union approximates F . The quality of the approximation strongly depends on the position and tangent continuity across panel boundaries: *Divergence* quantifies the spatial gap between adjacent panels, while the *kink angle* measures the jump in normal vectors across the panel intersection curves.

Figure 3: Terminology and variables used in our algorithm. The reference surface F and the initial curve network C are given as part of the design specification. The optimization solves for the mold depot $\mathcal{M}$, the panel-mold assignment function A , the shape parameters of the molds, the alignment transformations T_i , and the curve network samples c_k .

Curved panels are commonly produced using a manufacturing mold M_k . We call the collection $\mathcal{M} = \{M_1, \dots, M_m\}$ with $m \leq n$ the *mold depot*. To specify which mold is used to produce which panel(s), we define a panel-mold assignment function $A : [1, n] \rightarrow [1, m]$ that assigns to each panel index the corresponding mold index. The arrangement of panels in world coordinates is established by rigid transformations T_i that align each panel P_i to the reference surface F . Panels produced from the same mold are sub-patches of the mold surface and need not be congruent.

Let $c(M_k)$ be the fabrication cost of mold M_k and $c(M_k, P_i)$ the cost of producing panel P_i using mold M_k (see also Figure 9). The total cost of panel production can then be written as

$$\text{cost}(F, \mathcal{P}, \mathcal{M}, A) = \sum_{k=1}^m c(M_k) + \sum_{i=1}^n c(M_{A(i)}, P_i). \quad (1)$$

Ideally, the same mold will be used for the fabrication of multiple panels to reduce cost. The choice of panel types depends on the desired material and on the available manufacturing technology. Materials may be transparent or opaque, and include glass, glass-fibre reinforced concrete or gypsum, various types of metal, and wood.

Currently we support five panel types: planes, cylinders, paraboloids, torus patches, and general cubic patches. Planar panels are easiest to produce, but result in a faceted appearance when approximating curved freeform surfaces, which may not satisfy the aesthetic criteria of the design. A simple class of curved panels are cylinders, a special case of single-curved (developable) panels. Naturally, such panels can lead to a smooth appearance only if the given reference surface exhibits one low principal curvature. General freeform surfaces often require double-curved panels to achieve the desired tolerances in fitting error, divergence, and kink angles. We consider three instances of such panels: paraboloids, torus patches, and cubic patches. The former two carry families of congruent profiles (parabolas and circles, respectively), which typically simplifies mold production. Cubic panels are most expensive to manufacture, but offer the highest flexibility and approximation power. Thus a small number of such molds are often indispensable to achieve a reasonable quality-cost tradeoff.

Curve Network. The intersection curves between adjacent panels are essential for the visual appearance of many designs (see Figure 2) and typically affect the structural integrity of the building, as they often directly relate to the underlying support structure. An initial layout of these curves is usually provided by the architect as an integral part of the design. While small deviations to improve the paneling quality are typically acceptable, the final

solution should stay faithful to the initial curve layout and reproduce the given pattern as well as possible by the intersection lines of adjacent panels. Our paneling algorithm supports arbitrary curve network topology and is not restricted to predefined patterns. The collection of all panel boundary curves forms the *curve network* that we denote with $\mathcal{C}$. Projecting $\mathcal{C}$ onto the input freeform surface F yields a partitioning of F into a collection $\mathcal{S} = \{s_1, \dots, s_n\}$ of *segments* s_i . The panel P_i associated with segment s_i can be created by trimming the aligned mold surface $M_i^* := T_i(M_{A(i)})$. The trim curves are obtained by projecting the network curves associated with segment s_i onto M_i^* .

2.2 The Paneling Problem

We formulate the paneling problem as follows: Approximate a given freeform surface by a collection of panels of preferred types such that the total production cost is minimized, while the paneling respects pre-defined thresholds on divergence and kink angle between adjacent panels, and reproduces the initial curve network as well as possible. A closer look at this specification reveals that any solution of the paneling problem has to address the following central aspects:

- **Mold depot:** determine the number and types of molds that should be fabricated.
- **Assignments:** find the optimal assignment function to establish which panel is best produced by which mold.
- **Registration:** compute the optimal shape parameters for each mold and the optimal alignment of each panel such that the reference surface is faithfully approximated, thresholds on kink angles and divergence are met, and the panel intersections curves respect the design intent.

Mold depot and assignment function determine the total cost of fabrication, while registration affects the quality of the rationalization. Minimizing fabrication cost calls for a maximum amount of mold reuse and the wider use of panels that are geometrically simple and thus cheaper to manufacture (see Equation 1). On the other hand, achieving the design constraints on the paneling quality pushes the solution towards more complex panel shapes with less potential for mold reuse.

The high intricacy of the paneling problem arises both from the large scale of typical projects (1k – 10k panels) and the tight global coupling of objectives. Neighboring panels are strongly linked locally through kink angle and divergence measures, but also subject to a highly non-local coupling through the assignment function that

Figure 4: Our algorithm allows controlling the amount of deviation from the reference surface, shown here for the example of the Lilium tower. Larger deviations enable a more cost-effective solution using cheaper panels, while still satisfying the thresholds on kink angle and divergence.

facilitates mold reuse. The cost-quality tradeoffs involved in using different mold types add additional complexity: It is obvious that we want to use as many cheap, simple molds as possible. However, adding one expensive mold might save us from having to add many cheap molds whose total cost might be higher. Also, using a complex mold at certain places may enable the use of simpler panels in its surroundings.

3 Paneling Algorithm

A key design decision in our paneling algorithm is to represent the curve network explicitly as a set of polygonal curves, rather than computing the boundary curves from the intersection of neighboring panels (Figure 3). By integrating the corresponding curve vertices $\mathbf{c}_k$ as free variables in our system, we gain several important advantages. In particular, we

- avoid numerical instabilities when explicitly computing intersections of neighboring, nearly tangent-continuous patches,
- simplify the specification of surface fitting and continuity constraints across neighboring panels (see Section 3.1),
- achieve better quality at lower cost by allowing the curves to move away from the reference surface as part of the optimization (see Figure 4), and
- provide essential means of control for the designer, who can explicitly specify where neighboring panels should intersect.

An important aspect of the paneling problem is the need to simultaneously solve for both discrete and continuous variables. The discrete variables are the number and type of molds constituting the mold depot $\mathcal{M}$, and the panel-mold assignment function A . The continuous variables are the shape parameters of the molds, the transformations T_i that align mold $M_{A(i)}$ to segment s_i , and the positions of the curve network samples $\mathbf{c}_k$.

Given a mold depot $\mathcal{M}$ and assignments A , we use a continuous nonlinear least-squares optimization (Section 3.1) to globally improve the curve network and mold alignments. Such a least-squares optimization alone, however, does not ensure that the user-specified thresholds on kink angles and divergence are met, nor can it be used to determine a mold depot or assignments. The core challenge and main contribution of our algorithm is to find a mold depot and an as-

signment function that minimize cost and at the same time meet the specified thresholds. Minimizing cost in geometric terms means approximating the design with as many simple and repetitive elements as possible. In Section 3.2 we show how this difficult objective can be mapped to a generalized set cover problem, a classical problem in computer science. We present an efficient approximation algorithm that on its own provides a solution to the paneling problem. Due to the exponential complexity of possible mold-panel assignments this algorithm can only employ local registrations and therefore neither supports globally coupled continuous registration nor optimization of the curve network to allow deviations from the reference surface. Since the algorithm described in Section 3.2 enables solving for all the discrete variables we call it *Discrete Optimization* in contrast to the global *Continuous Optimization* of Section 3.1. To combine the strengths of the discrete and continuous optimization we present an iterative scheme in Section 3.3 to interleave the two to obtain a powerful global paneling algorithm. We demonstrate in Figure 8 how this interleaved iteration significantly improves the paneling compared to only applying the algorithm described in 3.2.

3.1 Continuous Optimization Step

The continuous step aims at reducing the deviation to the reference surface, the divergence, and the kink angles by optimizing the continuous variables, i.e., the shape parameters of the molds, the rigid alignments of mold surfaces to segments, and the positions of the curve network samples. During this optimization, the panel-mold assignments are kept fixed.

A mold surface M_k is specified in a canonical coordinate system as a function of a set of parameters that we store in a vector $\mathbf{m}_k$ (see Appendix). Depending on the type of mold surface, $\mathbf{m}_k$ contains zero (plane) to six entries (cubic). The rigid transformations T_i that align mold surface $M_{A(i)}$ to the corresponding surface segment s_i are initialized by the the local alignments computed in the discrete optimization. The curve network $\mathcal{C}$ is represented by samples $\mathbf{c}_l \in \mathbb{R}^3$, $l = 1, \dots, L$, where $L \approx 9|\mathcal{S}|$ in all our examples. We denote with $i(l)$ and $j(l)$ the indices of the two surface segments adjacent to the boundary curve sample $\mathbf{c}_l$. For notational brevity of the subsequent formulae, we omit the more complex adjacencies at corner points of the curve network. These are handled analogously by considering all combinations of neighbors.

The explicit representation of the curve network enables direct control of the panel intersection curves and allows formulating objective functions for inter-panel position and tangent continuity using simple closest-points projections. Effectively, the curve network serves as a “glue” that binds the panels to the surface at the curve samples. Since the manufacturing process mandates that the panel surfaces are smooth and exhibit a low variation of curvature, we found that additional samples in the interior of the segments are not required.

Surface fitting. The deviation of the curve network from the underlying reference surface F is measured as

$$E_{\text{fit}} = \sum_{l=1}^L \|\mathbf{c}_l - \mathbf{f}_l\|^2, \quad (2)$$

where $\mathbf{f}_l$ is the closest point on F from $\mathbf{c}_l$.

Divergence. We measure the divergence indirectly using the distances between the curve network samples and the closest points of the adjacent aligned mold surfaces as

$$E_{\text{div}} = \sum_{l=1}^L \|\mathbf{c}_l - \mathbf{x}_{i(l)}\|^2 + \|\mathbf{c}_l - \mathbf{x}_{j(l)}\|^2, \quad (3)$$

where $\mathbf{x}_{i(l)}$ is the point closest to $\mathbf{c}_l$ on the aligned mold surface $M_{i(l)}^*$. The point $\mathbf{x}_{j(l)}$ is defined analogously. Minimizing E_{div} ensures that neighboring panels fit together nicely without leaving undesirable gaps and that the curve network stays spatially close to the aligned mold surfaces.

Kink angles. Tangent continuity is achieved by minimizing the kink angle energy

$$E_{\text{kink}} = \sum_{l=1}^L \|\mathbf{n}(\mathbf{x}_{i(l)}) - \mathbf{n}(\mathbf{x}_{j(l)})\|^2, \quad (4)$$

where $\mathbf{n}(\mathbf{x}_{i(l)})$ is the normal vector of the aligned mold $M_{i(l)}^*$ at $\mathbf{x}_{i(l)}$, and analogously for $\mathbf{n}(\mathbf{x}_{j(l)})$ (Figure 3).

Curve fairness. The curve network is coupled to the surface through the surface fitting energy, but can move freely along the surface. To maintain the design intent encoded in the initial curve layout, we need to prevent strong undulations and large tangential motions of the boundary curves. Fortunately, the explicit representation of the curve network allows us to directly control its motion. We avoid tangential drift entirely by restricting the displacements of the curve network to the normal direction of the underlying reference surface. This simplification has the additional benefit of reducing the number of optimization variables for the curve network to one third and thus significantly improves performance. The shape of the curves is controlled using an additional fairness term on the curve deformation. Let $\mathbf{c}_l = \mathbf{c}_l' + d_l \mathbf{n}_l'$, where $\mathbf{c}_l'$ is the initial position of the curve network sample on the reference surface, $\mathbf{n}_l'$ is the corresponding surface normal, and d_l is the displacement magnitude. A fairness term can then be defined directly on the normal displacements as

$$E_{\text{fair}} = \sum_{(j_1, j_2) \in \mathcal{C}} (d_{j_1} - d_{j_2})^2, \quad (5)$$

where (j_1, j_2) is an index pair denoting an edge of the polygonal representation of the curve network $\mathcal{C}$.

Panel centering. With the exception of plane and cylinder molds, we need an additional energy term to facilitate mold reuse. Panels that are fabricated from the same mold should cover a similar region of the mold surface to reduce wastage in fabrication. For this purpose we add a centering energy that minimizes tangential drifting of the mold center away from the segment center:

$$E_{\text{cen}} = \sum_{i=1}^n \|\mathbf{b}_i - \mathbf{p}_i\|^2. \quad (6)$$

Here $\mathbf{b}_i$ is an approximation of the barycenter of segment s_i computed as the average of all adjacent curve segment samples, and $\mathbf{p}_i$ is the projection of $\mathbf{b}_i$ onto the mold-surface normal at the center of the aligned mold M_i^* . For torus molds we include an additional variable that positions the center along the defining circle (see Appendix). Figure 5 illustrates the effect of the centering energy.

Global optimization. The above continuous energy terms are combined into a global energy

$$E = \alpha_{\text{fit}} E_{\text{fit}} + \alpha_{\text{div}} E_{\text{div}} + \alpha_{\text{kink}} E_{\text{kink}} + \alpha_{\text{fair}} E_{\text{fair}} + \alpha_{\text{cen}} E_{\text{cen}}, \quad (7)$$

where the weights α allow additional control of the optimization. The unknown variables in E that we solve for are: the mold parameters $\mathbf{m}_k$, the rigid transformations T_i , and the positions of the curve network samples $\mathbf{c}_l$. For the T_i we use the formulation based on instantaneous velocities that ensures rigidity of the transformation [Pottmann and Wallner 2001]. The global energy E is minimized using a Gauss-Newton solver.

Figure 5: A low weight for E_{cen} may result in scattered molds (red) over a base primitive leading to wastage and increased cost. A higher weight results in tightly clustered molds (green) enabling cost-effective production of a significantly smaller mold.

Parameters. As detailed in Section 3.2, our system allows the user to provide thresholds ϵ for divergence and δ for kink angles. To balance the corresponding optimization objectives, we set $\alpha_{\text{kink}} = (\epsilon/\delta)^2 \alpha_{\text{div}}$ in Equation 7. This ensures that a divergence of ϵ receives the same penalty as a kink angle of δ . The free parameters have an intuitive meaning and offer direct control of the paneling quality. The specific values for these parameters depend on the user preferences. Figure 4 illustrates the effect of α_{fit} to control the deviation from the reference surface using values of 0.0001 and 1000, respectively. Figure 5 demonstrates the influence of α_{cen} . For all other examples we use the same set of parameters: $\alpha_{\text{fit}} = 1$, $\alpha_{\text{div}} = 1000$, $\alpha_{\text{fair}} = 1$, $\alpha_{\text{cen}} = 10$.

3.2 Discrete Optimization Step

The discrete optimization finds a mold depot and a corresponding panel-mold assignment function that minimize cost while respecting the specified divergence and kink angle thresholds. During the discrete optimization, the curve network is fixed. The essential geometric information required in this step is stored in the curve network samples and corresponding normal vectors that we obtain by averaging for each curve network sample the normals of all neighboring panels of the current solution. We therefore represent each surface segment by the associated curve network samples and normals, and use these to evaluate divergence and kink angles. This segment representation allows an efficient fitting of molds to segments based on local registration independent of the neighboring panels, using local versions of the divergence, kink angle, and centering energies (Equations 3, 4, 6).

To enrich the mold depot, we first create new candidate molds by locally aligning each of the available panel types to each of the segments. These $t|\mathcal{S}|$ new molds, where t is the number of panel types, together with the existing mold depot of the current solution form the set $\mathcal{M}'$ of candidate molds. The algorithm optimizes for a new mold depot as a subset $\mathcal{M} \subseteq \mathcal{M}'$ of all candidate molds that enables paneling at minimal cost, while satisfying the current thresholds on divergence and kink angle.

Set cover. Assume that we have determined for each mold $M_k \in \mathcal{M}'$ the set $\mathcal{S}_k = \{s_{k1}, \dots, s_{kl}\}$ of surface segments that can be approximated by M_k within the prescribed tolerances. We can compute the potential fabrication costs attributed to set $\mathcal{S}_k$ as $c(\mathcal{S}_k) = c(M_k) + |\mathcal{S}_k| c(M_k, P_*)$, i.e., the cost of the mold plus for all assigned panels the cost $c(M_k, P_*)$ of producing a panel with mold M_k .

Finding the optimal mold depot and mold-segment assignments now amounts to covering the set of segments $\mathcal{S} = \{s_1, \dots, s_n\}$ with sets $\mathcal{S}_k$ of minimal total cost. This optimization is reminiscent of the classical *weighted set cover* problem [Johnson 1974], where

the weights correspond to production costs. While this problem is known to be NP-hard, a polynomial-time approximation strategy can be used to find an approximate solution whose cost is guaranteed to be within $\log n$ times the cost of the optimal solution. It has been shown that $\log n$ is the best possible approximation ratio of any polynomial-time algorithm [Feige 1998].

What distinguishes our setting from the classical weighted set cover problem is that each segment is eventually assigned to only one set. Thus our weights (costs) depend on which segment is assigned to which set: Once a mold has been chosen, all segments covered by the corresponding set have to be removed from the sets of the remaining candidate molds. Nevertheless, the proof for the $\log n$ approximation ratio of the polynomial-time weighted set cover algorithm directly translates to our more general setup.

Algorithm. Let $S' \subseteq S$ be the set of all uncovered segments at any point of the algorithm. We define the *efficiency* of a set S_k relative to S' as the function $\phi(S_k, S') = |S_k|/c(S_k)$, where both the size of the set S_k and its cost are adapted to only consider elements in S' . Efficiency measures the number of panels for segments in S' that can be produced by mold M_k relative to the corresponding cost. Let σ be the unknown collection of covering sets and σ' the collection of all sets S_k that have not yet been chosen for σ .

Our algorithm to determine the covering sets starts with an empty collection σ . Thus the collection σ' contains the sets S_k of all candidate molds in $\mathcal{M}'$ and the set S' contains all segments that can be covered with the sets in σ' . We then successively add the set S_i with highest efficiency to the current solution. The segments of S_i are removed from S' , the efficiency of all remaining sets is updated, and the algorithm is iterated until all segments are covered.

```

 $\sigma \leftarrow \emptyset, \sigma' \leftarrow \{S_1, \dots, S_{|\mathcal{M}'|}\}, S' \leftarrow S_1 \cup \dots \cup S_{|\mathcal{M}'|}$ 
while  $S' \neq \emptyset$ 
    eval.  $\phi(S_k, S') \quad \forall S_k \in \sigma' \quad \text{update efficiencies}$ 
     $S_i \leftarrow \arg \max_{S_k \in \sigma'} \phi(S_k, S') \quad \text{set with max. efficiency}$ 
     $\sigma \leftarrow \sigma \cup \{S_i\} \quad \text{add to covering sets}$ 
     $S' \leftarrow S' - S_i, \sigma' \leftarrow \sigma' - \{S_i\} \quad \text{remove covered segments}$ 
end

```

The covering sets in the solution σ define the mold depot $\mathcal{M}$. For all segments that cannot be covered by any of the initial sets S_k we add and assign the best-fit cubic mold. Once the mold depot has been selected, a gather step re-computes the panel-mold assignment function by selecting for each segment the cheapest valid mold from the depot $\mathcal{M}$. In case a segment is covered by multiple molds of the same cost, we pick the one with the smallest maximal distance to the curve network.

Sets initialization. A critical bottleneck in the set cover optimization is the estimation of the initial sets S_k . To avoid the exhaustive computation of aligning every mold with every segment using nonlinear registration, we implement a pruning step that discards mold-segment pairs based on a conservative estimate of the corresponding registration distance. For this purpose, we introduce a 6D metric space that facilitates efficient approximate distance computations between molds and surface segments without the need for an explicit alignment.

We first select for each segment s_i its least-squares optimal cubic mold $C_i \in \mathcal{M}'$ as the segment's representative. The cubic mold C_i provides the best local approximation possible for this segment in the current configuration. To estimate whether a given mold sur-

Figure 6: Progressive panel assignments for different iterations (decreasing thresholds) of the paneling algorithm (Lilium tower).

face $M_k \in \mathcal{M}'$ is a suitable candidate for segment s_i , we compute an approximate alignment distance between M_k and C_i . By mapping both patches to points in a 6D space as described in the Appendix, this computation amounts to simple Euclidean distance evaluations. Finding potential molds to fit a given segment then only requires a range query in the 6D space that is performed efficiently using spatial data structures such as a kd-tree [Arya and Mount 1993]. This pruning step typically leads to a reduction in explicit mold-segment registrations to about 1% – 5%. For this substantially smaller set, we then perform the full nonlinear registration as described above to build the sets S_k .

3.3 Interleaved Iteration

Our algorithm iteratively executes the discrete and continuous optimization steps in an interleaved fashion. The exponential number of possible panel-mold assignments mandates the use of local registration methods in the discrete optimization. As a result, the set cover algorithm produces a mold depot and assignment function that is too costly for the given thresholds. Therefore, we start the iterations with thresholds well above the specified target values ϵ for divergence and δ for kink angles. Thresholds are then successively reduced after each iteration until the target values are reached.

We use the following fixed scheduling scheme for all our examples: We start the optimization with $\epsilon' = \epsilon + 10\text{mm}$ and $\delta' = \delta + 5^\circ$ and apply 10 iterations of alternating discrete/continuous optimization steps, reducing the thresholds in each iteration by 1mm and 0.5° .

The algorithm starts with a mold depot consisting of a single plane. We perform an initialization step that first determines all panels that do not meet the target thresholds (ϵ, δ) . These panels are replaced by locally fitting a separate mold of the cheapest type that satisfies the thresholds (ϵ', δ') . Subsequently, we perform a continuous optimization step. This initialization is applied before every discrete optimization to reverse unsuccessful assignments of the set cover and avoid a premature fixing of panel-mold assignments.

The inner loop of the optimization consists of the following steps:

- discrete optimization
- continuous optimization
- re-initialization
- reduce ϵ', δ'

At the end of the optimization we apply a discrete step using the target thresholds to generate the final mold depot and assignment function. Figure 6 illustrates the iterations of the paneling algorithm. Panels for which the target thresholds cannot be met mandate the fabrication of custom freeform molds. An important benefit of our approach is that the global discrete/continuous optimization leads to a small number of these custom molds (see also Figure 8).

Figure 7: Paneling results with varying kink angle thresholds δ and fixed divergence thresholds $\epsilon = 6\text{mm}$ for the design of the National Holding Headquarters. The images on the right show a solution using only planar panels of which 3,796 do not meet the prescribed divergence threshold. The zooms show reflection lines to illustrate inter-panel continuity which successively improves with lower kink angle thresholds.

Figure 8: Comparison of different methods for the same quality thresholds. State-of-the-art commercial tools only support a greedy panel assignment based on local fitting (top). Just one single application of our discrete optimization greatly reduces cost without loss in surface quality (middle). The full paneling algorithm interleaving discrete optimization with global continuous registration produces a high quality paneling (bottom). This solution contains 90% single curved panels and a very small number of custom molds, leading to a significantly reduced cost compared to greedy and local methods. The zoom on the right shows that our algorithm supports arbitrary curve network topology, including t-junctions. (Zaha Hadid Architects, Dongdaemun Design Plaza and Park, Seoul.)

Figure 9: Glass vs. concrete. Different relative costs for mold fabrication and panel production for two different materials affect the distribution of panel types. For glass, costs for producing double curved molds are significantly higher than for concrete, resulting in a solution with more cylindrical panels. (Zaha Hadid Architects, interior skin of Heydar Aliyev Merkezi Cultural Center, Baku.)

4 Evaluation and Discussion

Means of control. We allow the designer to explore the space of paneling solutions using two main modes of control: (i) specifying the quality using thresholds for divergence and kink angles (see Figure 7), and (ii) allowing the paneling solution to deviate from the reference surface in order to achieve a more cost-effective solution while maintaining the original design intent (see Figure 4). Due to our choice of α_{fit} (Section 3.1) the maximal deviation from the reference surface is within 10-20cm for all our results (average panel size $4m^2$). Unless stated otherwise, we use the cost set of glass panels as denoted in Figure 9 and the target thresholds $(\epsilon, \delta) = (6mm, 3^\circ)$.

Case Studies. Paneling solutions strongly depend on the design (reference surface and curve network), the choice of material, and the manufacturing technology. We explore the implications of these design decisions on various contemporary architectural freeform projects from leading architects in Figures 9, 10, and 11. Figure 8 demonstrates the effectiveness of the core components of our paneling algorithm. The most complex case study is the Dongdaemun Design Plaza and Park project (Figure 8) consisting of 8,385 panels. For this example, one step of discrete optimization takes on average 30 minutes, while one step of continuous optimization requires about 10 minutes, leading to a total computation time of 430 minutes. All computations are performed on a 3GHz Intel® Core™2 Duo with 3Gb of memory.

Figure 10: Paneling solutions for different curve networks on the Lilium tower. Significant changes in the curve network layout result in different paneling solutions. The left curve network, being closer to a conjugate network [Pottmann et al. 2007], is better suited for rationalization with planar and single curved panels.

Extensibility. Our framework for paneling architectural freeform surfaces takes important practical requirements into consideration. It is extensible in various ways to account for further specific needs and constraints. These include restrictions on mold reuse (restricting the use of the same mold to only parts of the surface determined by the assembling schedule; restricting the number of panels to be produced by the same mold), region-specific quality criteria (for example using higher thresholds for hardly visible parts), or explicit specification of panel types allowed for selected segments. Currently our system supports five important panel types. Incorporating further types is not difficult, but may require some thought to integrate them into a metric space for fast panel-segment distance computations (see Appendix). Since all our architectural designs are segmented rather uniformly, our current cost model does not take panel sizes into account. Panel sizes could be easily included, however, by specifying a cost per square meter for each panel type and adopting the discrete optimization accordingly.

Figure 11: Distribution of panel types and mold reuse for two different refinement levels of the curve network. The histograms on the right illustrate mold reuse for the top image, where the x -axis denotes how many panels can be produced by a mold and the y -axis indicates the number of panels corresponding to these molds. (Mario Bellini Architects, Rudy Ricciotti, Museum of Islamic Arts at Louvre Museum, Paris.)

Limitations. Although we achieve reasonable computation times for our examples, very large models ($> 50k$ segments) may only be effectively handled by splitting them into several parts that are treated separately. Even if splitting of a large model is not necessary, the performance may not allow the user to fully exploit the potential of the available control mechanisms (thresholds, allowed deviation from reference surface).

In our current problem formulation we expect an initial curve network as part of the design. Certain materials like gypsum are suitable for producing freeform surfaces without visible seams. Hence the aesthetics of the curve network plays a minor role. While our framework could be extended to allow stronger movement of the curve network, the additional flexibility of freely optimizing over the curve layout is currently not exploited in our algorithm.

Future Work. There are a number of desirable extensions to our method that constitute challenging problems for future research. One direction would be to incorporate further important aspects into the optimization, e.g., structural feasibility and efficiency of the underlying support structure, especially if it is aligned with the network of panel boundaries. This may be done using properly simplified mechanical models, in extension of work by Whiting et al. [2009], and should lead to a new powerful tool for form-finding. The optimization could also include energy performance values or try to optimally integrate solar panels. These and related extensions towards the optimized use of the building would be in the spirit of the Building Information Modeling paradigm. To our knowledge the paneling problem is also of interest to the ship-hull construction industry. We plan to investigate this application in future work.

5 Conclusion

We have cast the paneling of architectural freeform surfaces into a global optimization problem that takes into account panel production cost, reuse of manufacturing molds, and various constraints on surface quality. The solution is based on a combination of discrete and continuous optimization, as well as a new inter-panel distance approximation for complexity reduction. We have demonstrated the performance of our system on several complex freeform designs by leading contemporary architects.

Although our research in paneling is primarily inspired by architecture, the problem is of relevance in the broader context of geometric modeling. Approximation of a general shape by a nearly smooth union of aesthetically arranged panels of a limited number of types is a significant extension of the state-of-the-art in surface rationalization. Simultaneously deciding on the number of different panel types required, while optimizing their parameters leads to an interesting and unusual mixture of discrete and continuous global optimization with potential applications in other domains. Finally, mold reuse by making relations explicit across different object parts and removing redundancies in the shape information content, goes beyond symmetry detection and symmetrization towards global shape understanding.

Acknowledgements. We would like to thank Yves Brise, Peter Kaufmann, Sebastian Martin, and Heinz Schmiedhofer for their help with the project. Special thanks to Zaha Hadid Architects and Waagner Biro for providing the architectural datasets. This work was supported by the Swiss National Science Foundation, the Austrian Research Promotion Agency grant number 824197, and the European Community's 7th Framework Programme under grant agreement 230520 (ARC). Niloy Mitra was partially supported by a Microsoft outstanding young faculty fellowship.

References

- ARYA, S., AND MOUNT, D. M. 1993. Approximate nearest neighbor queries in fixed dimensions. In *SODA*, 271–280.
- BOBENKO, A., AND SURIS, YU. 2009. *Discrete differential geometry: Integrable Structure*. Graduate Studies in Math. AMS.
- COHEN-STEINER, D., ALLIEZ, P., AND DESBRUN, M. 2004. Variational shape approximation. *ACM Trans. Graphics* 23, 3, 905–914.
- FEIGE, U. 1998. A threshold of $\ln n$ for approximating set cover. *JACM* 45, 4, 634–652.
- FU, C.-W., LAI, C.-F., HE, Y., AND COHEN-OR, D. 2010. K-set tilable surfaces. *ACM Trans. Graphics* 29, 3.
- GLYMPH, J., SHELDEN, D., CECCATO, C., MUSSEL, J., AND SCHÖBER, H. 2002. A parametric strategy for freeform glass structures using quadrilateral planar facets. In *Acadia*, 303–321.
- GOLOVINSKIY, A., PODOLAK, J., AND FUNKHOUSER, T. 2009. Symmetry-aware mesh processing. *Mathematics of Surfaces*.
- JOHNSON, D. S. 1974. Approximation algorithms for combinatorial problems. *Journal of Computer and System Sciences* 9, 256.
- JULIUS, D., KRAEVOY, V., AND SHEFFER, A. 2005. D-charts: Quasi-developable mesh segmentation. *Computer Graphics Forum (Proc. Eurographics)* 24, 3, 581–590.
- LI, B., SCHNABEL, R., SHIYAO, J., AND KLEIN, R. 2009. Variational surface approximation and model selection. *Computer Graphics Forum (Proc. Pacific Graphics)* 28, 7, 1985–1994.
- LIU, Y., POTTMANN, H., WALLNER, J., YANG, Y.-L., AND WANG, W. 2006. Geometric modeling with conical meshes and developable surfaces. *ACM Trans. Graphics* 25, 3, 681–689.
- MITRA, N. J., GUIBAS, L. J., AND PAULY, M. 2006. Partial and approximate symmetry detection for 3D geometry. *ACM Trans. Graphics* 25, 3, 560–568.
- MITRA, N. J., GUIBAS, L. J., AND PAULY, M. 2007. Symmetrization. *ACM Trans. Graphics* 26, 3, 63.

- PAULY, M., MITRA, N. J., WALLNER, J., POTTMAN, H., AND GUIBAS, L. J. 2008. Discovering structural regularity in 3D geometry. *ACM Trans. Graphics* 27, 3, 43.
- PODOLAK, J., SHILANE, P., GOLOVINSKIY, A., RUSINKIEWICZ, S., AND FUNKHOUSER, T. 2006. A planar-reflective symmetry transform for 3D shapes. *ACM Trans. Graphics* 25, 3, 549–559.
- POTTMANN, H., AND WALLNER, J. 2001. *Computational Line Geometry*. Springer-Verlag New York, Inc.
- POTTMANN, H., LIU, Y., WALLNER, J., BOBENKO, A., AND WANG, W. 2007. Geometry of multi-layer freeform structures for architecture. *ACM Trans. Graphics* 26, 3, 65.
- POTTMANN, H., HOFER, M., AND KILIAN, A., Eds. 2008. *Advances in Architectural Geometry*. Vienna.
- POTTMANN, H., SCHIFTNER, A., BO, P., SCHMIEDHOFER, H., WANG, W., BALDASSINI, N., AND WALLNER, J. 2008. Freeform surfaces from single curved panels. *ACM Trans. Graphics* 27, 3, 76.
- SHAMIR, A. 2008. A survey on mesh segmentation techniques. In *Computer Graphics Forum*, vol. 27, 1539–1556.
- SHELDEN, D. 2002. *Digital surface representation and the constructibility of Gehry's architecture*. PhD thesis, M.I.T.
- SINGH, M., AND SCHAEFER, S. 2010. Triangle surfaces with discrete equivalence classes. *ACM Trans. Graphics* 29, 3.
- VARADY, T., AND MARTIN, R. 2002. Reverse engineering. In *Handbook of CAGD*, 651–681.
- WHITING, E., OCHSENDORF, J., AND DURAND, F. 2009. Procedural modeling of structurally-sound masonry buildings. *ACM Trans. Graphics* 28, 5, 112.
- WU, J., AND KOBELT, L. 2005. Structure recovery via hybrid variational surface approximation. *Computer Graphics Forum (Proc. Eurographics)* 24, 3, 277–284.
- YAN, D.-M., LIU, Y., AND WANG, W. 2006. Quadric surface extraction by variational shape approximation. In *GMP*, 73–86.

Appendix

We efficiently estimate registration errors in a 6D metric space to avoid more than 95% of the costly nonlinear mold-segment alignments for the discrete optimization step of Section 3.2. We define a mapping of planes, paraboloids, and cubic patches into this space and show how cylinders and tori can be approximated with cubics. Representing segments by their best fitting cubic polynomial then allows calculating upper bounds on the registration error through simple 6D Euclidean distances computations.

Distances between cubic polynomials. We represent a cubic polynomial patch in the principle frame of the patch center as

$$P_i(x, y) = a_i x^2 + b_i y^2 + c_i x^3 + d_i x^2 y + e_i x y^2 + f_i y^3.$$

Our goal is to find potential mold candidates for a given segment. We thus consider the domain $D = [-L, L]^2$, where L is half the segment side length computed on the initial curve network. Since these approximate distances are only used for pruning, it is sufficient to assume a quadratic shape even for non-quadratic segments. We define the L_2 -distance $E_{ij} = E(P_i, P_j)$ between two polynomials P_i and P_j by optimizing over the relative shift in z -direction

$$E_{ij} = \min_{\tau} \iint_D (ax^2 + by^2 + cx^3 + dx^2 y + exy^2 + fy^3 + \tau)^2 dx dy,$$

Figure 12: Maximal kink angle and divergence averaged by fitting 400 randomly selected molds to all the 7265 segments of the National Head Quarters design. The values are sorted by our approximate panel-segment distances to demonstrate the strong correlation of the approximative and the exact fitting errors.

where $a := a_i - a_j$, etc. Since the integral can be factored into the form $A\tau^2 + 2B\tau + C$, we obtain $E_{ij} = C - B^2/A$ for the optimal value of $\tau = -B/A$. Substituting $\bar{d}_i = d_i + f_i$ and $\bar{e}_i = c_i + e_i$ and reordering terms yields

$$E_{ij} = 4L^2 \left(\frac{4L^4}{45} a^2 + \frac{4L^4}{45} b^2 + \frac{L^6}{15} \bar{d}^2 + \frac{L^6}{15} \bar{e}^2 + \frac{8L^6}{105} c^2 + \frac{8L^6}{105} f^2 \right).$$

By setting $\bar{E} = \sqrt{E/(4L^2)}$, our distance measure between cubic patches corresponds to the canonical Euclidean distance where every cubic is represented as a point in the 6D Euclidean space as

$$\mathbf{P}_i = \left(\frac{2L^2}{3\sqrt{5}} a_i, \frac{2L^2}{3\sqrt{5}} b_i, \frac{L^3}{\sqrt{15}} \bar{d}_i, \frac{L^3}{\sqrt{15}} \bar{e}_i, \frac{\sqrt{8}L^3}{\sqrt{105}} c_i, \frac{\sqrt{8}L^3}{\sqrt{105}} f_i \right)^T.$$

Since we are interested in the shape of the cubic rather than its specific orientation, we consider for every cubic P_i all four right handed coordinate transformations and another four with a flipped normal direction z :

$$P_i^* = \{P_i(x, y), P_i(-y, x), P_i(-x, -y), P_i(y, -x), \\ -P_i(x, -y), -P_i(y, x), -P_i(-x, y), -P_i(-y, -x)\}$$

which correspond to eight points in the derived Euclidean space. We finally define our approximative distance measure as

$$d(P_i, P_j) = \min_{P'_j \in P_j^*} \bar{E}(P_i, P'_j),$$

which defines a metric on the space of cubic polynomials.

Cylinder and torus. Planes and paraboloids can be directly described as cubic polynomials with zero higher order coefficients. For cylinders and tori, we use the Taylor expansion expressed in the coordinate system of the principle frame at a center point on the mold as an approximation. For cylinders with radius R the only nonzero coefficient is thus $a_i = 1/(2R)$.

Torus molds are defined by a meridian radius R_1 , a parallel radius R_2 , and angle σ that determines where on the meridian the mold center lies. Since this representation has three parameters, we use the third-order Taylor expansion

$$T_3(x, y) = -\frac{1}{2} \left(\frac{1}{R_1} x^2 + \frac{\cos \sigma}{R} y^2 + \frac{R_2 \sin \sigma}{R^2 R_1} xy^2 \right), \quad (8)$$

where $R := R_2 + R_1 \cos \sigma$, and $-1/R_1$ and $-(\cos \sigma)/R$ are the principal curvatures at the mold center.