Procedural Mesh Features applied to Subdivision Surfaces using Graph Grammars

W. Thaller¹¹, U. H. Augsdörfer²¹ and D. W. Fellner³ ^{1,2} ¹Institut für ComputerGraphik & Wissensvisualisierung, TU Graz, Austria ²TU Darmstadt & Fraunhofer IGD, Germany

Abstract

A typical industrial design modeling scenario involves defining the overall shape of a product followed by adding detail features. Procedural features are well-established in computer aided design (CAD) involving regular forms, but are less applicable to free-form modeling involving subdivision surfaces. Current approaches do not generate sparse subdivision control meshes as output, which is why free-form features are manually modeled into subdivision control meshes by domain experts. Domain experts change the local topology of the subdivision control mesh to incorporate features into the surface, without increasing the mesh density unnecessarily and carefully avoiding the appearance of artefacts.

In this paper we show how to translate this expert knowledge to grammar rules. The rules may then be invoked in an interactive system to automatically apply features to subdivision surfaces.

Keywords: Procedural features; feature-based modeling; free-form surface features; graph grammars; Catmull-Clark subdivision surfaces; automating changes in mesh topology

1 1. Introduction

The design of products that have to fulfil engineering requirements as well as aesthetic criteria, typically involves freeform shapes. The styling activity of aesthetic product design comprises two steps [1]: First the product's overall shape is defined. This is followed by local refinements where features are added to the overall shape. Often, the same or similar features are repeatedly applied to one design.

A procedural feature consists of a set of parameters along with an algorithm for applying the feature to an underlying model. The designer can then manipulate the feature directly on a high semantic level of abstraction. The underlying shape scan be edited while leaving the procedural features in place. Also, the designer is able to control specific aspects of a feature shape, while the overall shape remains fixed. This is referred to as *feature-based modeling*.

For feature-based modeling to work, features must be wellte defined in terms of their parameters. Therefore, feature-based modeling was first introduced in the context of solid modeling, where procedural features have become firmly established. A cl cylindrical hole drilled into an overall shape is the classic example of a procedural feature. The use of free-form surface features, rather than regular-shaped features, is an active area of research [2].

Because feature-based modeling systems need to be compatible with a larger *modeling pipeline*, it is desirable for the
system's input and output to be a standard CAD representation.
Non-uniform rational B-splines (NURBS), a patch-based
surface representation, are the current standard for free-form
modeling in CAD. Adding features to free-form surfaces represented by NURBS frequently involves increasing the resolution

³² of surface patches, which leads to many redundant vertices in ³³ the representation.

Subdivision surfaces, already established in the animation industry, have recently gained popularity as an alternative to NURBS in CAD. Subdivision methods are a generalisation of traditional spline patch methods to arbitrary topology; for example, Catmull-Clark [3] generalises bi-cubic patches.

Because subdivision surfaces may also include extraordinary vertices, that is vertices with a valency either more or less than the regular valence, features may be introduced into the surface by locally changing the topology of the control mesh.

However, changes to the topology of the control mesh may give rise to artefacts in the limit surface. To apply free-form design features to a subdivision surface, a CAD expert meticulously adjusts the resolution and the mesh topology of the subdivision control mesh in order to keep changes to the overall shape small and to avoid or hide the visual appearance of artefacts [4]. Once a good topology change is identified, the expert applies the same procedure each time the same or a similar feature has to be applied.

In this paper we describe how this expert knowledge can be formulated as graph grammar rules. We propose to use a graph grammar on top of a scripting language; the basic local modifications are scripted, and the graph grammar allows us to organize the problem.

57 2. Related Work

Grammars have proved very useful for the procedural mod-⁵⁹ eling of buildings where contex-free split grammars have been ⁶⁰ used [5, 6]. There, too, a grammar provides the organizing er paradigm for something that could theoretically also be hard- 113 3. Adding Features to Subdivision Surfaces 62 coded.

Procedural free-form features have been a lively area of re-63 64 search since the 1990s. For a comprehensive introduction to 65 and overview of feature-based modeling in CAD, we refer the 66 reader to Pernot et al. [2].

67 2.1. B-spline based methods

In [7], Pernot et al. introduce a method for using three-68 69 dimensional curves for defining deformations applied to B-spline ⁷⁰ surfaces. This method was later extended [8, 9, 10] as a method 71 for defining a range of freeform surface features on freeform 72 surfaces. Application of the features is done using a force-based 73 deformation system. The method can operate on NURBS data, 74 but no new control points are added. The input geometry must 75 have sufficient resolution to represent the features.

Chen et al. [11] solve this limitation by using hierarchical 77 NURBS [12] to represent the resulting surface. In this approach, 78 higher-resolution patches are used to refine the lower-resolution 79 input patches and add high-frequency detail. The resulting ge-80 ometry is tied to the use of non-standard hierarchical NURBS ⁸¹ as a geometry representation and can therefore not be used in 82 existing modeling pipelines.

Another way to avoid the problem of insufficient resolu-83 ⁸⁴ tion in the input control mesh is to generate the feature during 85 rendering instead of attempting to output a control mesh. Dis-⁸⁶ placement mapping [13] is a wide-spread and efficient method 87 for applying high-frequency detail to surfaces when the output ⁸⁸ is only used for visualization.

⁸⁹ 2.2. Subdivision based methods

A popular method for adding a limited repertoire of fea-90 91 tures to subdivision surfaces has been defined by De Rose et ⁹² al. [14]. In their approach, edges of a subdivision control mesh ⁹³ may be marked as being exempt from the subdivision process. ⁹⁴ This introduces sharp and semi-sharp creases into an otherwise 95 smooth limit surface. However, the output surface is not a gen-⁹⁶ eralisation of traditional spline patch methods, which typically ⁹⁷ is a requirement further down the product design pipeline.

Khodakovsky and Schröder [15] describe an algorithm al-⁹⁹ lowing the creation and manipulation of fine scale feature curves 100 on subdivision surfaces. Creation of the features happens dur-¹⁰¹ ing the subdivision process, so there is no sparse control net for 102 the resulting surface.

In the context of Sketch based modeling, Olsen et al. [16] 103 ¹⁰⁴ have developed a method for incorporating linear features into 105 subdivision control nets by locally increasing the mesh resolution. The transition between the higher-resolution patched faces and the lower-resolution surroundings is handled by a set of 107 ¹⁰⁸ fixed patching template. However, this approach may increase control mesh density more than necessary. 109

In product design today, topolgical changes to a subdivision 111 control net to incorporate high frequency features are still mod-112 eled manually by the CAD expert.

The problem of applying free-form features to a subdivi-114 115 sion surface can be seen as consisting of three subproblems, 116 namely 1) defining the feature, 2) changing the topology of the ¹¹⁷ subdivision control mesh, and 3) shifting the control points ap-118 propriately.

Of these, the second step is the most time consuming and 119 120 the one where expert knowledge is required.

Typically, a CAD expert with a thorough understanding of 122 the surface representation carefully designs the subdivision con-123 trol mesh in order to locally increase the resolution around the 124 area where the feature is to be placed. This is not a trivial task:

CAD modelers typically take great care to avoid the appear-126 ance of surface artefacts.

In regular regions surface artefacts are known to arise if fea-127 128 tures are not aligned but run skew to grid lines of the control 129 mesh [4, 17, 18]. Because of this CAD modelers take great 130 care to align features with the underlying grid.

Locally refining a subdivision control mesh to incorporate 132 the feature typically gives rise to irregular regions: For a subdi-133 vision scheme based on quadrilateral meshes these occur around ¹³⁴ vertices with more or less than four edges, referred to as extraor-135 dinary vertices, and non-quad faces, which give rise to extraor-¹³⁶ dinary vertices after one subdivision step.

Subdivision surfaces do not guarantee C^2 continuity at ex-138 traordinary vertices [19], and undesirable artefacts are likely to 139 appear [4, 20, 21].

Typically, when applying features to a subdivision surface, 141 expert mesh modelers meticulously identify changes to the topol-142 ogy of subdivision control meshes such that the visibility of un-143 avoidable artefacts around irregular regions is minimized.

Furthermore, when different features meet in one place on 145 a surface, they have to interact. The interaction of features is 146 domain-specific and cannot be determined without access to 147 domain-specific expert knowledge.

As is clear from the above, incorporating features to a free-149 form design is time consuming and often requires an expert un-150 derstanding of the underlying geometric representation in order 151 to integrate features.

Once a method for good integration of a feature to a subdivi-153 sion control mesh has been identified, the expert CAD designer 154 has to use the same methods repeatedly in order to manually 155 incorporate free-form features of similar type into subdivision 156 surface control meshes. There are currently insufficient intelli-157 gent tools to support or automate this laborious task.

Automating some of the repetitive tasks of the feature-based 158 159 subdivision modeling will accelerate the design process and 160 frees the designer from needing extensive knowledge of the un-¹⁶¹ derlying geometry representation.

162 4. Contribution of the Paper

We demonstrate that grammar rules are suitable to translate 163 ¹⁶⁴ expert knowledge of how to change the topology of a control ¹⁶⁵ mesh to incorporate a feature. Compared to existing methods ¹⁶⁶ for automatically applying features to subdivision meshes, the ¹⁶⁷ rule-based approach adds significantly fewer control points to 168 the mesh.

Changing the topology of a control mesh unavoidably intro-169 170 duces changes to the limit surface. To keep the alteration of the 171 original surface to a minimum we apply an optimization step ¹⁷² after the changes in mesh topology.

To demonstrate our approach, we have developed a proto-173 ¹⁷⁴ type system to automatically change the topology of a Catmull-175 Clark subdivision control mesh [3] to apply various surface fea-176 tures with only a minor increase in resolution to achieve good 177 visual properties.

178 5. Our Approach

A mesh data structure can be interpreted as a graph. There-179 $_{180}$ fore, graph grammars can be used to define operations on meshes. $_{223}$ We employ graph grammars to express the expert knowledge of 182 how to integrate features into a subdivision control mesh in a way that changes its mesh topology with respect to pre-defined 183 requirements. 184

The vertices and halfedges of the subdivision control mesh 18 ¹⁸⁶ form the nodes of the graph. Arbitrary attributes can be added to 187 the nodes. Relationships between these objects are represented 188 as edges in the graph.

Features are defined in terms of rules that find the position 189 of a given pattern in the graph and then apply some local pre-190 ¹⁹¹ defined operation to the mesh at that position. These rules capture the CAD expert's knowledge of the required changes in the 192 topology of the subdivision control mesh to integrate a feature. 193 The feature may then be applied using a simple point-and-194 click interface by non-expert users. 195

Our prototype system is built on top of the Generative Mod-196 ¹⁹⁷ eling Language (GML) [22], a scripting language with special ¹⁹⁸ support for the manipulation of Catmull-Clark [3] subdivision ²³⁷ surfaces. Local modifications to the mesh are described using 199 the scripting language; the questions of where and when the 200 local modifications are to be applied, and of how they interact 201 with each other, are handled by the graph grammar. While any 202 scripting language with library support for manipulating subdi-203 vision control meshes could be used, using a GML script has 205 two advantages: First, GML's use of Euler Operators as mesh 206 manipulation primitives always guarantees that the invariants of the mesh data structure are preserved. Second, the description is often shortened because unchanged parts of the graph need 209 not be repeated.

210 5.1. Meshes and Graphs

Using a standard halfedge datastructure, a mesh consists of 21 vertices, halfedges and faces. Each halfedge connects exactly 212 213 two vertices, its source and its target, and belongs to exactly one face. Each halfedge is related to exactly one other halfedge ²¹⁵ by a symmetric relation. These two halfedges connect the same ²¹⁶ two vertices, but in opposite directions, i.e. source and target 217 are swapped. For each vertex, all outgoing halfedges form a ²¹⁸ cycle, and all halfedges belonging to a face form a cycle. As ²¹⁹ an example, Figure 1 shows a two-dimensional representation



Figure 1: A halfedge mesh representing a tetrahedron. Vertices A, B, C and D are connected with pairs of halfedges in opposing directions.

²²⁰ of the halfedge mesh of a tetrahedron. Meshes with a boundary 221 can be handled by representing the boundary as a single, non-²²² planar face that is marked as invisible.

Figure 2 shows a tetrahedron represented as a graph. To ob-224 tain a graph representation of a mesh, we choose to map each 225 element of the mesh, i.e. each vertex and each halfedge, to a 226 node in the graph, and the relationships between the elements 227 to edges in the graph. Furthermore, grammar rules are allowed 228 to add arbitrary additional labeled edges to the graph, and to 229 assign arbitrary attributes to nodes (name-value pairs). Geo-230 metric information, such as the position of a vertex, is stored as 231 an attribute for each vertex.

For the purpose of our system, we can define:

Definition 1. A graph G is a tuple (N, E, A, L, V) where

• N is a set of nodes

234

235

236

239

- $E \subseteq N \times N \times L$ is a set of edges
- L is a set of labels
- V is a set of possible attribute values
- $A \in N \times L \rightarrow V$ is a function defining attributes for each node.

The set V of attribute values corresponds to the set of all 240 ²⁴¹ values supported by the scripting language. The set L is the set 242 of identifiers.

Definition 1 defines a directed graph whose edges are unique 243 for each label, i.e. given a pair of nodes $a, b \in N$, there is at most $_{245}$ one edge from *a* to *b* for each edge label.

246 5.2. Graph Grammars

Graph grammars, by definition, consist of a set of rules ap-247 ²⁴⁸ plied to a graph [23]. Each rule matches a subgraph against the 249 graph; this subgraph is then replaced by a different subgraph as ²⁵⁰ defined in the rules. Rules are applied until no rule matches any 251 more.

As an example let us look at the simple rule that replaces ²⁵³ one quad with two triangles, as seen in Figure 3. In this example 254 the rule defines what elements (nodes and edges) of the graph 255 on the left side correspond to which elements of the graph on 256 the right side.



Figure 2: The graph corresponding to the mesh of a tetrahedron. The circled nodes *A*, *B*, *C* and *D* correspond to the vertices of the tetrahedron. The two-letter nodes correspond to the halfedges, i.e. node *AB* represents the halfedge leading from *A* to *B*, and node *BA* represents its "mate", going from *B* to *A*. The various graph edges represent different relationships between vertices and halfedges. Blue and green denote relations of halfedges to their source and target vertices, respectively. Red edges denote a symmetric relation between halfedges. The dotted edges represent faces for which associated halfedges form a circle. Thus, the edges *AB* – *BD* – *DA* form a cycle corresponding to the face *ABD*.

In our system, the left sides of the rules are represented usmatched, and the right sides are represented as a script that describes the operation performed to change the subgraph matched by the left side. This approach is analogous to the use of operations in split grammar systems such as CGA shape [6].

²⁶³ Our example rule from Figure 3 can be textually represented ²⁶⁴ as follows in our system:

```
ule Example {
265
       e1(color="red")
266
            faceCCW> e2
                                      e3
                          faceCCW>
267
                          faceCCW>
268
            faceCCW>
                      e4
                                      e1
269
       action { e1 e3 makeEF pop }
270
271
```

272 The "left side" of the rule — the code between the first pair 273 of curly braces above — represents a graph pattern. We match 274 against four nodes in the graph, e1 through e4, which are re-275 lated to each other by faceCCW edges, which correspond to the 276 dotted edges in Figure 2. Additionally, to match our pattern, 277 the first of these nodes must be marked by a color="red" 278 attribute. The "right side" of the rule consists of GML code 279 that defines what should happen whenever the left side of the 280 rule is matched: e3 e1 makeEF pop. Because GML is a 281 stack-based language, this means that the makeEF operation 282 is applied to two parameters, e1 and e3 (which were defined 283 while matching the left side of the rule), and its result is dis-



Figure 3: Example of a graph grammar rule (top) and its application to a mesh (bottom): Find a quad-face, made up of vertices A, B, C, and D connected via halfedges e1, e2, e3 and e4, where one edge (e1) is marked with the "red" attribute, as shown on the left. Split this quad-face into two triangles along the AC diagonal. The system matches the pattern to the mesh (lower left) and applies the rule everywhere it matches, yielding a new mesh (lower right).

²⁸⁴ carded (pop). The makeEF Euler operator (supplied by GML) ²⁸⁵ creates a new face (F) by adding a new edge (E) between its ²⁸⁶ parameters.

We can also script arbitrary additional conditions to the pattern to formulate any special cases used by the expert. In the above example, we could decide to make the rule apply to all quads, but with the orientation of the diagonal split chosen in such a way that the newly inserted diagonal is the shorter diagonal.

For more complex rules, where new vertices are inserted, the scripting language code calculates their coordinates. In practice, this is often a simple linear combination of the existing vertices.

This is not the first time grammar concepts have been applied to meshes. Spicher et al. [24] describe a system based on topological collections rather than graph grammars, intended for the declarative specification of subdivision algorithms. By contrast with our system, the right sides of the rules in [24] are defined using declarative expressions that evaluate to topological collections of cells (i.e., collections of vertices, edges and faces). While the declarative approach leads to a more mathematically beautiful formulation, the fact that cells on the righthand side have to be explicitly generated by this expression may lead to increased verbosity.

308 5.3. Surface Optimization

³⁰⁹ It is important that the application of a feature keeps the ³¹⁰ alteration of the original overall shape to a minimum.

311 troduces changes to the surface. After performing the topol-³¹³ ogy changes, but before shifting the control points to create ³¹⁴ the feature, we therefore add an optimization step to minimize ₃₆₆ by applying pre-defined rules at locations in the control mesh 315 the changes to the limit surface of the overall shape. Introducing this optimization step allows many graph grammar rules to ³¹⁷ make the simplifying assumption that the surface is flat; and ³¹⁸ new control points need only be placed correctly within the flat ³⁷⁰ types of surface features can be created from very simple rules. ³¹⁹ face of the control mesh. Errors introduced by the assumption of flatness are corrected by the optimization. 320

We have chosen a simple optimization algorithm that yields 32 322 reasonable results very quickly. Essentially, for each control point we pick a fixed direction and move the control point along that direction until its limit point falls on the desired surface. 324

Below, L^* is the original limit surface before any rules were 325 ³²⁶ applied, c_i^0 are the control points after application of the graph grammar rules, and c_i^k are the control points after iteration k of 327 328 the optimization algorithm.

- 1. Determine the normal vectors n_i of the control mesh at 329 each control point c_i^0 . 330
- 2. Let l_i^0 be the limit points corresponding to the control 33 points c_i^0 . 332
- 3. Determine scalars s_i and points $p_i^0 = l_i^0 + s_i n_i$ such that ³⁸⁶ 333 $p_i^0 \in L^*$. 334
- 4. Let m_i be the normal vector of L^* at point p_i^0 . 335
- 5. Repeat for k = 0, 1, 2, ...: 336
- 6. Calculate the limit points l_i^k corresponding to the 337 control points c_i^k . 338

7. Determine scalars t_i^k and points $p_i^k = l_i^k + t_i^k m_i$ such 339 that $p_i^k \in L^*$. 340

8. Let
$$c_i^{k+1} = c_i^k + (p_i^k - l_i^k)$$
.

8. Let $c_i^{\kappa+1} = c_i^{\kappa} + (p_i^{\kappa} - l_i^{\kappa})$ 9. Terminate when $|t_i^k| < \varepsilon$. 342

For any of the meshes used in this paper the optimisation 343 ³⁴⁴ required less than twenty iterations.

Minimizing the error only for the limit points of the con-345 346 trol points, as opposed to over the entire surface between the 347 control points, makes this optimization algorithm very cheap, but also bears the risk of causing oscillation in the optimized 348 surface. However, the effect is strongest in the area where it is dominated by the displacement introduced by the feature itself, ³⁵¹ and thus mostly hidden from view.

352 6. Results

³⁵⁴ have built a prototype system to apply free-form surface features to a Catmull-Clark subdivision surface. 355

356 358 click, edges in a subdivision control net. The system allows 415 More elaborate markings may create features with intersections $_{359}$ the user to either mark edges to automatically add the feature at $_{416}$ and T-junctions as in Figure 6(e). the centre of the edges or the user may choose specific locations 417 360 $_{362}$ is meant to cross. These markings are represented as attributes $_{419}$ 6(e). Without the optimisation step (left), the error introduced

Changing the topology of a control mesh inadvertently in- 363 according to Definition 1. The designer can pre-select different ³⁶⁴ types of features, which leads to different attributes being set. ³⁶⁵ The feature is then incorporated to the subdivision control mesh 367 where these attributes are found. Rules are applied until no 368 further rules match.

> In this section, we will demonstrate how several standard 369 ³⁷¹ Figure 4 shows a set of diagrams that describe a simple rule 372 set for creating linear surface features. For brevity's sake, we ³⁷³ only show the graphical representation of the rules, and we omit 374 trivial variations of the depicted rules.

> Figure 5 shows the application of a groove-feature on a sub-375 ³⁷⁶ division surface. It is modeled by choosing the type of feature 377 to be applied ("groove") and then marking edges in the con-378 trol mesh via mouse click with the attribute "red". The system 379 then automatically inserts three new vertices along each marked ³⁸⁰ edge in the control mesh according to Rule #1 shown in Figure ³⁸¹ 4. The topology changes are the same, independent of the pa-³⁸² rameters of the feature, i.e. the width or depth of the groove.

> When the marked edges are on opposite sides of a face in the 384 control mesh, as in Figure 5, the feature is applied like a sim-³⁸⁵ ple translational sweep. This involves #2 from Figure 4. Using the graph grammar approach, an arbitrary number of different 387 scenarios of feature application may be formalised, handling 388 various situations, such as corners (Rule #3), crossings (Rule 389 #4) and three-way junctions (Rules #5 and #6). T-junctions in ³⁹⁰ subdivision surfaces have to be handled very carefully. The in-³⁹¹ serted control mesh lines cannot simply end at a T-junction, as 392 this would lead to irregular faces that markedly degrade sur-³⁹³ face quality due to the appearance of artefacts. The patterns 394 seen in Rules #5 and #6 are often used to minimize these prob-395 lems. Observe how the right sides of these rules contain only ³⁹⁶ quadrilateral faces. Rule #6 also moves the geometric position ³⁹⁷ of the vertex originally at the top right of the pattern to be in 398 line with the feature that ends in the T-junction. Just like with ³⁹⁹ the new vertices added by the rules, the scripting language code 400 can calculate the coordinates using the simplifying assumption 401 that the surface is flat; the placement with respect to the third ⁴⁰² dimension is determined by optimization (section 5.3).

403 Figures 6(b) to 6(e) show the results of applying features ⁴⁰⁴ to the Catmull-Clark subdivision surface shown in Figure 6(a). 405 The three columns show on the left the control mesh modi-406 fied to incorporate the feature, in the centre the limit Catmull-407 Clark surface with the feature and on the right the limit sur-408 face colourised to highlight differences between the new and 409 the original limit surface. By extending the rule database with To demonstrate the approach presented in this paper we 410 more rules, different feature types can be applied, such as a 411 groove feature (Figure 6(b)), a ridge feature (Figure 6(c)) or 412 the two-dimensional "dent" feature shown in Figure 6(d). Sur-In our prototype interface, the location where the feature $_{413}$ faces 6(b) to 6(d) are applied by choosing the type of feature is to be applied is defined by the user by marking, via mouse $_{414}$ followed by marking the control mesh as shown in Figure 6(a).

Figure 7 illustrates the effect of the optimization step dealong an edge in the subdivision control mesh which the feature 418 scribed in section 5.3 applied to the feature applied in Figure



Figure 4: A rule set for generating simple seams, grooves or ridges as surface features. Initially, all mesh edges to be crossed by the seam are marked in red by the user. Rule #1 splits those edges into smaller edges by adding three vertices along each marked edge. The remaining rules connect the newly created points in different situations: straight feature across a quad (Rule #2), corners (Rule #3), crossings (Rule #4), tee-crossings (Rule #5 and #6). Two extra graph edges, shown as blue dashed arrows, are added by Rule #1 to communicate to the later rules which mesh edges belong together. When none of the rules apply anymore, the actual feature is created by shifting all vertices marked in green along the surface normal vector towards the interior of the surface for simple seams and grooves, or away from the surface for ridges.



Figure 5: Topology changes to incorporate a groove-feature into a subdivision surface by applying graph grammar rules shown in Figure 4: The designer marks the affected edges in the control mesh (a). The system automatically applies Rule #1 to each marked edge: Marked edges are split by inserting additional control points (b). The resulting control mesh (b) contains several 10-gons that match the left side of Rule #2. Rule #2 is applied, that is new points are connected across the faces (c). Finally, the centre vertices of the seams are shifted along the negative direction of their normal vectors, resulting in the final control mesh (d), the limit surface of which includes the groove feature.

⁴²⁰ by applying the feature is more pronounced, especially in ar-⁴²¹ eas where the underlying surface has a higher curvature. With ⁴²² the optimisation, the error is not eliminated but significantly re-⁴²³ duced (right).

An experienced designer will align features with grid lines. Area feature does not follow the existing control lines, artefacts will appear [4, 17, 18]. Typically, CAD experts will avoid this situation in practice by planning ahead when designing the verall shape of a model.

Our method is, however, not restricted to aligning features 429 430 with grid lines. It can be used to also formulate changes to 431 the topology to also apply features skew to the grid lines. In ⁴³² Figure 8 we used the system presented in this paper to apply 433 features not aligned with the control mesh. The three columns 434 in the figure show again the modified control mesh (left), the 435 limit surface (centre) and the error of the new surface to the 436 original saddle shape without the feature (right). In Figure 8(a) 437 a free-shape feature intersects a feature aligned with the grid 438 lines. In Figure 8(b) the boundaries of the dent feature follow 439 an arbitrary curve. As long as the feature to be added keeps 440 clear of existing control points, the existing rules for straight 441 (Rule #2) and corner (Rule #3) features are sufficient. For the 442 case where the feature comes too close to an existing vertex, 443 we can add special-case rules, Rules #7 and #8 (Figure 9), that 444 shift the vertex out of the way. This introduces additional error 445 into the surface, but it allows the feature to be realised without 446 adding many new vertices and without non-local changes to the control mesh. 447

It is worth noting that the graph grammar concept is capable of expressing different approaches to the problem of adding for eatures to the mesh; for example, patch patterns for adaptive subdivision described in [16] map directly to grammar rules (Figure 10). These rules serve as a general mechanism for losides cally increasing the resolution of a mesh in order to add hightion detail to an otherwise sparse control mesh. By com-



Figure 7: Effect of the optimization step described in Section 5.3. Without optimization (a), the error is much more pronounced than with optimization (b).



Figure 8: Non-grid-aligned features, created using two additional rules (Figure 9) that move aside the control points of the underlying surface where they would otherwise interfere with the features.



Figure 6: After marking the affected edges on the original mesh (a), various features can be automatically applied: grooves (b) and ridges (c), as well as two-dimensional features like an indentation (d). Given different markings, features may also interact with each other or form T-junctions and crossings as in (e).

Figure 9: Two rules for enabling non-grid-aligned features by moving aside control points that are in the way. As before, a half-edge marked red indicates that a feature will cross this edge; the exact place where the feature should cross is stored as a further attribute. These rules are conditional — they are only applied if the feature would otherwise be too close to the vertex in the center of the pattern. The amount by which the central vertex is shifted is likewise determined by script code.



Figure 10: Adaptive control mesh refinement using the patching patterns from [16] can be directly implemented using graph grammar rules to serve as a general mechanism for adding detail. Compared to specialized rules, this usually leads to a greater increase in resolution.



Figure 11: Complex features may be defined using graph grammars: A piped seam is a seam where an extra piece of material, the *piping*, is sewn into the seam. For our purposes, we can model this by using a variant of Rule #1 that applies a different profile to the edges.

⁴⁵⁵ parison, a special-purpose ruleset as described in this paper uses 456 far fewer control points to model a feature.

457 7. Example Application

Let us now consider the real-world application of a system 458 459 which supports a non-expert designer to incorporate different 460 types of seam features to the upholstery of chairs modeled using ⁴⁶¹ Catmull-Clark subdivision surfaces (Figure 13).

The most important attribute for defining a seam feature is 462 463 its profile, where the user may adjust parameters such as width and depth of the seam. 464

We demonstrate the ability of our system to deal with in-465 teractions between different features by using different profiles 466 to create different kinds of seams, such as piped seams, where 468 an extra piece of material, the so-called piping, is sewn into the 469 seam. This gives the seam a characteristic appearance with a ⁴⁷⁰ protruding center. We can simulate this using a more complex 471 profile for the seam, as shown in Figure 11. Figures 13(d) and 472 (e) show a surface feature pattern involving piped seams applied 503 8. Discussion 473 to our example model.

When two features cross it is sometimes sufficient to apply 474 475 existing rules to the mesh. In other situations new mesh patterns 476 are required to avoid artefacts caused by the interaction of both 477 features. The ability to formulate ad-hoc special cases used by experts as rules that can be applied automatically is one of the 478 479 strengths of the approach presented in this paper.

Interactions between piped seams and regular seams are 480 such special cases; a three-way junction where a regular seam 481 482 ends in a piped seam (Figure 13(d)) is very similar to the a regu-483 lar three-way junction, but a crossing between a piped seam and





Figure 12: The crossing between two features (a regular seam and a piped seam) may lead to artefacts if standard crossing rules (Rule #4, Figure 4) are applied (a). To avoid artefacts special rule may be employed to handle these situations. In this example, visually more pleasing results are achieved if the regular seam ends where it meets the piped seam and starts again on the other side (b).

484 a regular seam feature, as shown in Figures 13(e), leads to arte-485 facts in the surface if handled by a variant of Rule #4. Figure 486 12(a) shows a piped seam crossing a regular seam, constructed 487 using the regular crossing rule, Rule #4. Artefacts occur in the 488 limit surface where the many parallel edges running cause a 489 "kink" in the piping, see Figure 12(a) on the right. To improve 490 the appearance of the subdivision surface we add another rule 491 to the graph grammar to handle this special case in a more ap-⁴⁹² propriate manner, such as to use the strategy from Rule #5 to 493 eliminate the extra control points. The result is shown in Figure 12(b), where no artefacts are visible in the limit surface.

Note that the case of two piped seams crossing each other 496 is yet another special case. It is not a trivial crossing, since ⁴⁹⁷ the pipings for the two seams cannot occupy the same space. ⁴⁹⁸ One of the pipings needs to end where it meets the other seam, 499 and resume on the other side. This special case can also be ⁵⁰⁰ automated by our system; it cannot be dealt with by a general ⁵⁰¹ algorithm, as it is a special case that arises purely from the ap-502 plication domain.

The methods that experts use to manually add a feature to a ⁵⁰⁵ mesh depend on the existing structure of the mesh, and so does ⁵⁰⁶ the graph-grammar based approach.

Some of the other methods described in Section 2 apply ⁵⁰⁸ brute force to this problem: While the grid lines of the control ⁵⁰⁹ mesh do not follow the generated feature, thus decreasing the 510 quality of the surface due to artefacts, the resulting errors can ⁵¹¹ be made arbitrarily small by increasing the resolution [17].

This is a general trade-off. Generating a higher-resolution 513 output reduces the amount of artefacts, while at the same time 514 makes the output less useful further down the product design



Figure 13: Various different patterns of seam features applied to the same model of a chair. Figures (a), (b) and (c) showcase the crossing rule (Rule #4), the corner rule (Rule #3) and a T-junction rule (Rule #6) when applying regular seams. Figures (d) and (e) demonstrate the application of different feature types, namely regular and piped seams. While the T-junction in (d) follows the same rule as involved in (c), the crossing in Figure (e) requires a special case rule to improve the appearance of the limit surface. Finally, (f) shows that features need not be aligned to grid lines.

⁵¹⁵ pipelines, since alterations to the dense control mesh are more ⁵⁶⁹ 516 difficult to realise.

In our approach the output is a modified subdivision control 518 mesh which incorporates a chosen surface feature at a small in-⁵¹⁹ crease in control mesh density which is only locally applied. 520 Therefore, the resulting control mesh is ideally suited for fur-521 ther manual modification.

At this point, writing or extending a graph grammar rule-522 523 set is not a trivial task and cannot be left to the designer; while the graph grammar abstraction markedly reduces the effort re-525 quired compared to hard-coding the mesh manipulations as a 526 monolithic script or program, an expert is still required. Apply-527 ing surface features by invoking the rules, by contrast, is now ⁵²⁸ open to non-experts. Once a rule has been formulated it may be applied countless times via a mouse click. 529

Special-purpose rules are not to be understood as rules that 530 are used only once, but rather as a way of creating domainspecific modeling tools. 532

If no rule is available for a specific situation, a modeling 533 ⁵³⁴ system can signal the failure to the user by marking that specific 535 part of the model. The designer can then deal with this particu-⁵³⁶ lar special case by hand while still profiting from the automated 537 feature application everywhere else in the model. Manual postmodification is made possible by the fact that the output is a 538 subdivision control mesh that is still as sparse as possible.

Graph rewriting is Turing complete, so concerns about the 540 ⁵⁴¹ undecidability of the halting problem are reasonable at a first 542 glance. We have, however, found that this is no more a con-⁵⁴³ cern than the Turing completeness of the underlying scripting 544 language used for defining the local mesh modifications. Al-⁵⁴⁵ most all rules that were used for the figures in this paper are 546 non-recursive, with the exception of Figure 10, which shows ⁵⁴⁷ a recursive refinement process that can easily be controlled by ⁵⁴⁸ including the desired refinement level as a condition.

549 9. Conclusion

We have shown that graph grammars can be employed to 550 551 capture expert knowledge about what mesh operations are nec-552 essary to incorporate surface features into subdivision control ⁵⁵³ meshes. Compared to other methods to automatically apply sur-⁵⁵⁴ face features, this method produces an output mesh that has significantly fewer redundant control points and thus better quality 555 control meshes. 556

Using this method to automate the application of surface 55 558 features into existing subdivision surface based modeling work-559 flows will significantly speed up the modeling process. Adoption in practice is likely to happen more quickly than for other ⁵⁶¹ automated approaches as the output mesh is not only more ef-562 ficient in terms of the number of control points used, but also very close to the result of manual modeling, allowing for a mixand-match approach during the adoption phase.

565 References

- [1] Fontana M, Giannini F, Meirana M. A free form feature taxonomy. In: 566
- 567 Computer Graphics Forum; vol. 18. Wiley Online Library; 1999, p. 107-18. 568

[2] Pernot JP, Falcidieno B, Giannini F, Léon JC. Incorporating free-form features in aesthetic and engineering product design: State-of-the-art report. Computers in Industry 2008;59(6):626-37. doi:10.1016/j.compind.2008.03.004.

570

571

572

574

575

576

577

578

579

580

581

582

583

584

585

587

588

589 590

591

592

593

594

595

596

- 573 [3] Catmull E, Clark J. Recursively generated b-spline surfaces on arbitrary topological meshes. Computer-Aided Design 1978;10(6):350-5. doi:10.1016/0010-4485(78)90110-0.
 - Sabin M, Barthe L. Artifacts in recursive subdivision surfaces. In: Cohen [4] A, Merrien JL, Schumaker LL, editors. Curve and Surface Fitting: St Malo 2002. Nashboro Press, Brentwood, TN; 2003, p. 353-62.
 - Wonka P. Wimmer M. Sillion F. Ribarsky W. [5] Instant architecture. ACM Transactions on Graphics 2003;22(3):669-77. doi:10.1145/882262.882324
 - Müller P, Wonka P, Haegler S, Ulmer A, Van Gool L. [6] Procedural modeling of buildings. ACM Trans Graph 2006;25(3):614-23. doi:10.1145/1141911.1141931.
 - Pernot JP, Guillet S, Leon JC. A shape deformation tool to model char-[71 acter lines in the early design phases. In: Shape Modeling International, 2002. Proceedings. 2002, p. 165-72. doi:10.1109/SMI.2002.1003542.
 - [8] Pernot JP, Falcidieno B, Giannini F, Guillet S, Léon JC. Modelling free-form surfaces using a feature-based approach. In: Proceedings of the Eighth ACM Symposium on Solid Modeling and Applications. SM '03; New York, NY, USA: ACM. ISBN 1-58113-706-0; 2003, p. 270-3. doi:10.1145/781606.781649.
 - Pernot JP, Falcidieno B, Guillet S, Leon JC. Aesthetic design of shapes [9] using fully free form deformation features. In: Tools and Methods for Competitive Engineering (TMCE); vol. 1. 2004, p. 143-54. ISBN 90-5966-023-4
- 597 [10] Cheutet V, Catalano C, Pernot J, Falcidieno B, Giannini F, Leon 3d sketching for aesthetic design using fully free-form de-598 T Computers & Graphics 2005;29(6):916 -30. 599 formation features. doi:http://dx.doi.org/10.1016/j.cag.2005.09.009. 600
- 601 [11] Chen W, Cai Y, Zheng J. Freeform-based form feature modeling using a hierarchical & multi-resolution nurbs method. In: Proceedings of the 602 9th ACM SIGGRAPH Conference on Virtual-Reality Continuum and Its 603 Applications in Industry. VRCAI '10; New York, NY, USA: ACM. ISBN 978-1-4503-0459-7; 2010, p. 179-84. doi:10.1145/1900179.1900218. 605
- 606 [12] Chen G, Esch G, Wonka P, Müller P, Zhang E. Interactive procedural street modeling. In: ACM SIGGRAPH 2008 papers. SIGGRAPH '08; New York, NY, USA: ACM. ISBN 978-1-4503-0112-1; 2008, p. 103:1-608 103:10. doi:10.1145/1399504.1360702. 609
- 610 [13] Cook LR. Shade trees. In: SIGGRAPH Proceedings. ACM Press; 1984, p. 223-231. 611
- 612 [14] DeRose T, Kass M, Truong T. Subdivision surfaces in character animation. In: Proceedings of the 25th annual conference on Computer graphics 613 and interactive techniques. SIGGRAPH '98; New York, NY, USA: ACM. 614 ISBN 0-89791-999-8; 1998, p. 85-94. doi:10.1145/280814.280826. 615
- Khodakovsky A, Schröder P. Fine level feature editing for subdivision 616 [15] surfaces. In: Proceedings of the Fifth ACM Symposium on Solid Mod-617 eling and Applications. SMA '99; New York, NY, USA: ACM. ISBN 618 1-58113-080-5; 1999, p. 203-11. doi:10.1145/304012.304033. 619
- 620 [16] Olsen L, Samavati FF, Sousa MC, Jorge JA. Sketch-Based Mesh Augmentation. In: Jorge JAP, Igarashi T, editors. Eurographics Workshop on 621 Sketch-Based Interfaces and Modeling. The Eurographics Association. 622 ISBN 3-905673-30-4; 2005, doi:10.2312/SBM/SBM05/043-052.
- 624 [17] Augsdörfer U, Dodgson N, Sabin M. Artifact analysis on b-splines, boxsplines and other surfaces defined by quadrilateral polyhedra. Computer 625 Aided Geometric Design 2011;28:177-197. 626
- 627 [18] Augsdörfer U, Dodgson N, Sabin M. Artifact analysis on triangular boxsplines and subdivision surfaces defined by triangular polyhedra. Com-628 puter Aided Geometric Design 2011;28:198-211. 629
- 630 [19] Peters J. Reif U. Shape characterization of subdivision surfaces — basic principles. Computer Aided Geometric Design 2004;21(6):585-99. 631
- 632 [20] Karciauskas K, Peters J, Reif U. Shape characterization of subdivision surfaces - case studies. Computer Aided Geometric Design 633 634 2004;21(6):601-14.
- 635 [21] Augsdörfer U. Dodgson N. Sabin M. Tuning subdivision by minimising Gaussian curvature variation near extraordinary vertices. 636 Computer Graphics Forum 2006;25(3):263-72. doi:10.1111/j.1467-8659.2006.00945.x. 638
- 639 [22] Havemann S. Generative Mesh Modeling. Ph.D. thesis; Institute of Com-

- puter Graphics, Faculty of Computer Science, Braunschweig Technical 640 University, Germany; 2005. doi:www.digibib.tu-bs.de/?docid=00000008; 641
- dOI: http://www.digibib.tu-bs.de/?docid=00000008. 642
- Rozenberg G, editor. Handbook of Graph Grammars and Computing by 643 [23] Graph Transformations, Volume 1: Foundations. World Scientific; 1997. 644
- ISBN 9810228848. 645 646 [24] Spicher A, Michel O, Giavitto JL. Declarative mesh subdivision using topological rewriting in mgs. In: Ehrig H, Rensink A, Rozenberg G, 647
- Schürr A, editors. Graph Transformations: 5th International Conference, 648
- ICGT 2010, Enschede, The Netherlands, September 27-October 2, 2010. 649
- Proceedings. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-642-15928-2; 2010, p. 298–313. doi:10.1007/978-3-642-15928-2_20. 650
- 651