

Model Synthesis: A General Procedural Modeling Algorithm

Paul Merrell and Dinesh Manocha
University of North Carolina at Chapel Hill

Abstract—We present a method for procedurally modeling general complex 3D shapes. Our approach can automatically generate complex models of buildings, man-made structures, or urban datasets in a few minutes based on user-defined inputs. The algorithm attempts to generate complex 3D models that resemble a user-defined input model and that satisfy various dimensional, geometric, and algebraic constraints to control the shape. These constraints are used to capture the intent of the user and generate shapes that look more natural. We also describe efficient techniques to handle complex shapes, highlight its performance on many different types of models. We compare model synthesis algorithms with other procedural modeling techniques, discuss the advantages of different approaches, and describe as close connection between model synthesis and context-sensitive grammars.

Index Terms—model synthesis, procedural modeling

1 INTRODUCTION

CREATING 3D digital content for computer games, movies, and virtual environments is an important and challenging problem. It is difficult because objects in the real-world are complex and have widely varying shapes and styles. Consider the problem of generating a realistic 3D model of an outdoor scene. Different applications may require many different types of models such as buildings, oil platforms, spacecrafts, roller coasters, and other man-made structures. Overall, geometric modeling is a creative and artistic process. However, current modeling systems can be cumbersome and the users often spend a lot of time performing routine and tedious tasks instead of making creative decisions.

Despite extensive work in geometric modeling for over four decades, it remains a time-consuming task. Learning how to use current geometric modeling tools can require significant training and even when the tools are mastered creating complex models is still difficult. With state of the art 3D CAD and modeling tools, the user can create simple geometric primitives and modify them using various transformations and geometric operations. Modeling complex environments such as cities or a landscapes requires the creation and manipulation of large numbers of primitives and can take many hours or days [31].

Many objects and environments contain repetitive and self-similar structures which can be modeled more easily using procedural modeling techniques. Procedural modeling techniques are designed to automatically or semi-automatically generate complex models. These include techniques based on shape grammars, scripting languages, L-systems, fractals, solid texturing, etc. These approaches have been used to generate many complex shapes, but each method is mainly limited to a specific class of models or requires considerable user input or

guidance.

In this paper, we address the problem of generating complex models using model synthesis. Model synthesis is a simple technique [26], [28] proposed to automatically generate complex shapes. The model synthesis algorithm accepts a simple 3D shape as an input and then generates a larger and more complex model that resembles the input in terms of its shape and local features. An example of this is shown in Figure 1.

Different procedural modeling techniques require varying degrees of user input. Using a high degree of user input has both advantages and disadvantages. Without sufficient user input, the result generated by a procedural modeling method may be too random and some parts of the generated 3D model may turn out to be different from the user’s original intent. With too much user input, the time required to adjust and manipulate the model could overwhelm the user. Ideally, the user could choose to provide any amount of input and the algorithm should be able to adjust accordingly. The user input can often be specified in the form of a set of constraints on the output. Any output that satisfies all of the user’s constraint is acceptable. Prior work in model synthesis [28] uses a minimal amount of user input in the form of a single adjacency constraint and may not give the user enough control over the result.

We present a novel model synthesis algorithm which enables the user to specify geometric constraints that give the user greater control over the results. We use dimensional, incident, algebraic, and connectivity constraints that have been used in CAD/CAM, geometric modeling, and robotics. The constraints are specified between a set of geometric objects and their features. These include spatial and logical constraints such as incidence, tangency, perpendicularity, and metric constraints such as distance, angle, etc. We use these constraints to capture the user’s intent, to prevent objects from becoming



Fig. 1. (a) From an example model specified by the user, (b) a model of several oil platforms is generated automatically by our algorithm. The shape of the output resembles the input and fits several dimensional and connectivity constraints. The height of the platform and the length and width of the beams are constrained to have a particular size. The shapes are constrained to be in four connected groups. Our algorithm can generate the new model in about half a minute.

unnaturally large or small, to generate more complex shapes, and to manage the objects' spatial distribution.

In order to satisfy the constraints, we represent local neighborhoods of the objects using Boolean expressions. The Boolean expressions are used to compute how different vertices, edges and faces of the synthesized model connect together. Furthermore, we present a scheme to incorporate dimensional and algebraic constraints into our model synthesis algorithm.

Like most procedural modeling techniques, our algorithm is primarily designed to work on objects that are self-similar. We demonstrate its ability to generate models of buildings, man-made structures, city streets, plumbing, etc.

A preliminary version of this paper appeared in [29]. We also compare model synthesis to other procedural modeling techniques which are often based on using grammars. We show that model synthesis is particularly useful for generating architectural shapes, but has difficulty with some curved and organic shapes. We also establish a close relationship between context-sensitive grammars and model synthesis by showing that a problem in one domain can be reduced to a problem in the other domain.

The rest of the paper is organized as follows. We give a brief survey of prior work on procedural modeling and geometric constraint systems in Section 2. Section 3 gives a brief overview of model synthesis and the constraints used by our algorithm. The overall constraint-based algorithm is described in Section 4 and we highlight its performance in Section 5. We analyze its performance and discuss its limitations in Section 6. We compare it to related work on procedural modeling in Section 7.

2 RELATED WORK

In this section, we briefly survey related work on procedural modeling, texture synthesis, and model synthesis.

2.1 Procedural Modeling

Many procedural modeling techniques have been developed over the last few decades. These techniques are very diverse, but most of them are targeted towards modeling a specific type of object or environment. Techniques based on fractal geometry achieved success modeling natural landscapes [13], [23], [32].

There is a long history of modeling plants procedurally. Many plant modeling techniques use a formal grammar called an L-system. L-systems were proposed by Lindenmayer as a general framework for describing plant growth and plant models [22], [36]. An L-system is a parallel rewriting system. L-systems can be extended to consider how plants interact with their environment as they grow [25]. Other techniques use sketches [8], photographs [38] or positional information [37] to influence the shape of the plant models.

Many procedural techniques are designed specifically for modeling urban models [45], [43]. Like many plant modeling techniques, some urban modeling techniques use L-systems. L-systems have been used to generate road networks and buildings on land between the roads [34]. Other grammars have been introduced specifically for modeling architecture. Shape grammars were introduced as a tool for analyzing and designing architecture [40], [12]. Wonka et al. introduced a related group of grammars called split grammars [49]. Split grammars operate by splitting shapes into smaller components and

can generate highly detailed models of architecture. Split grammars were further developed by Müller et al. [31]. They developed shape operations for mass modeling and for aligning many parts of a building’s design together.

Some techniques focus more on the 2D layouts of cities than on the 3D shapes of the buildings. Chen et al. [7] allow the users to edit a city’s street layout interactively using tensor fields. Aliaga et al. [2] generate street layouts using an example-based method. A related area of research is urban simulation which seeks to understand how various factors influence a cities development and growth over time [41], [44]. Aspects of urban simulation have been used in procedural modeling to produce more realistic models of cities [46].

Other techniques are designed to model smaller structures. Legakis et al. [21] propose a method for automatically embellishing 3D surfaces with various cellular textures including bricks, stones and tiles. Cutler et al. [9] developed a method for modeling layered, solid models with an internal structure. Another method has been developed to model truss structures by optimizing the locations and strengths of beams and joints [39]. Pottmann et al. [35] have developed algorithms based on Discrete Differential Geometry that determine how to arrange beams and glass panels so they form in the shape of a given freeform surface and satisfy various geometric and physical constraints.

Another way to model objects is to combine together parts of existing models interactively [14]. In this method, the user can search through a large database of 3D models to find a desired part, then cut the part out from the model, and stitch various parts together to create a new object.

2.2 Modeling with Constraints

There is rich literature in solid modeling on designing shapes that satisfy various geometric, parametric or variational constraints [6], [1]. There is also considerable work on solving geometric constrained systems and some excellent surveys are available [16], [18]. Geometric constraints are widely used in computer aided engineering applications [15] and also arise in many geometric modeling contexts such as virtual reality, robotics, molecular modeling, computer vision, etc. These constraints are used to incorporate relationships between geometric entities and features and thereby capture the intent of the designers. Our formulation of various constraints is similar, though our approach to satisfy these constraints during model synthesis is different. Besides geometric constraints, silhouette-based constraints are also used to model freeform objects using sketch-based interfaces [17], [33].

2.3 Texture Synthesis and Model Synthesis

The model synthesis algorithm itself has much in common with texture synthesis. The field of texture synthesis

has seen a proliferation of new algorithms and new ideas over the past decade. This section gives an overview of the most relevant developments and explains their relationship to model synthesis. A more comprehensive survey is given in [47].

Many texture synthesis algorithms were influenced by a seminal paper written by Efros and Leung [11]. Their algorithm is remarkably simple. It generates textures by adding pixels individually by finding a neighborhood that matches the neighborhood around the insertion point. There are several ways to accelerate texture synthesis such as by adding the pixels in a particular order [48], by adding patches of texture rather than individual pixels [10], [20], or by exploiting spatial coherence [3].

Texture synthesis has also been used to generate texture maps directly onto curved surfaces [42]. Texture synthesis has been extended into three dimensions is to create solid textures [19]. Texture synthesis has also been used to generate geometric textures [4], [51] which combine elements of texture mapping and modeling. They are used like texture maps to apply patterns to objects, but the patterns change the shape of the object to create effects like bumps or dimples or chain mail.

Texture synthesis was the inspiration behind model synthesis. Model synthesis was initially proposed by Merrell [26] and later extended to handle non-axis-aligned objects [28]. Both model and texture synthesis are designed to take a small sample as an input example and generate a larger result that resembles the input example.

Model synthesis relies on finding symmetric patterns. Many methods have been developed to find patterns in 3D models [24]. Bokeloh et al. [5] automatically identify symmetries within objects and use these symmetries to cut objects into pieces. By editing these pieces, a shape grammar is derived.

3 ALGORITHM

In this section, we give a brief overview of model synthesis and the constraints used in the algorithm.

3.1 Notation

Points and vectors are written in bold face, $\mathbf{x} \in \mathbb{R}^3$. Lower-case letters not in bold face are generally used to denote scalar variables, but there are a few exceptions. The variable h is used to denote the set of points within a half-space. The upper-case letters, E and M are used to denote the models. The model E is the input example model provided by the user. The model M is the new model generated by the algorithm. Each model is a set of closed polyhedra. The models E and M and the half-spaces h_i are represented in two different ways. A half-space h_1 could be represented as a set of points h_1 or as the characteristic function of that set $h_1(\mathbf{x})$ where $h_1(\mathbf{x}) = 1$ if $\mathbf{x} \in h_1$, otherwise $h_1(\mathbf{x}) = 0$. The complement of the half-space h_1 is written as either the set h_1^C or the function $\neg h_1(\mathbf{x})$.

3.2 Background

Our algorithm builds upon earlier work in model synthesis. In this section, we give a brief overview of a previous model synthesis algorithm [28]. The user provides an example model as the input. The example model is a set of polygons that form closed polyhedral objects. Model synthesis generates a new model M that resembles the example model E . In earlier work, it was assumed that the input was a single object, but we allow multiple objects in E . Let n be the number of different objects in E . We consider the example model to be a function $E(\mathbf{x})$ of a point in space \mathbf{x} where $E(\mathbf{x}) = k$ if \mathbf{x} is inside an object of type k where $1 \leq k \leq n$. If \mathbf{x} is not inside any of the objects, then $E(\mathbf{x}) = 0$. The function $M(\mathbf{x})$ is similarly defined for the new model M .

In the prior model synthesis algorithm, the output model only needed to satisfy a single constraint called the adjacency constraint. The adjacency constraint is defined on neighborhoods. A neighborhood around a point is just a set of points near it. In Figure 2, the neighborhoods surrounding the points \mathbf{a} , \mathbf{b} , \mathbf{c} , \mathbf{d} , and \mathbf{e} exactly match the neighborhoods surrounding the points \mathbf{a}' , \mathbf{b}' , \mathbf{c}' , \mathbf{d}' , and \mathbf{e}' . The neighborhood around a point \mathbf{x} matches the neighborhood around the point \mathbf{x}' if there exists $\epsilon > 0$ such that for all vectors δ where $\|\delta\| < \epsilon$

$$M(\mathbf{x} + \delta) = E(\mathbf{x}' + \delta). \quad (1)$$

The adjacency constraint states that for every point \mathbf{x} there exists a point \mathbf{x}' whose neighborhood matches according to equation 1. This constraint ensures that every neighborhood of M is the same as a neighborhood of E . A similar constraint is used in [5].

Figure 2 gives an overview of our approach. Starting with the input example shape $E(\mathbf{x})$ shown in Figure 2(b), the algorithm creates sets of parallel lines (or parallel planes in 3D) as shown in Figure 2(c). The output shape is generated on these sets of parallel lines. One possible output shape is shown in Figure 2(d). The lines (or planes in 3D) intersect at vertices.

Each vertex has a set of acceptable neighborhoods that match the input according to Equation 1. The vertex could be outside $M(\mathbf{x})$, inside it, or on its boundary which could be an edge, a face, or a vertex of the output shape. Each possible neighborhood is represented by a different possible state. One state might be a neighborhood which is on a face. A neighborhood on a face that has a different normal would be a different state. There are other states for neighborhoods on edges or vertices. Every neighborhood that is different according to Equation 1 is a different possible state. Several states are shown in Figures 3 and 4. Two states can be at adjacent vertices in Figure 2 if they have similar features along the direction in which they connect. For example, Figure 4 shows three states that could be beneath a particular state because they all share a vertical magnet edge that can connect the two states. Adjacent states

which do not share common features conflict because they cannot connect together.

After creating the planes, the next part of the algorithm is to assign states to each vertex without assigning two adjacent states that conflict. We keep track of a list of every possible state that could be assigned to each edge and each vertex. This list is long initially, but it shortens as we assign more states. Each assigned state is associated with a set of states that could be adjacent to it. Neighboring states outside this set conflict with the assigned state and get removed from the list. This removal may, in turn, expose other conflicting states which are also removed. This process is repeated until no more states need to be removed. We continue to assign states to each vertices and then update the list of possible states until every vertex has been assigned a single state.

3.3 Geometric Constraints

Our approach uses several geometric constraints to capture the user’s intent and to control the shape of the synthesized model. To describe different constraints, we borrow terminology from the solid modeling and CAD literature [1]. Figure 2(a) provides an overview of how the various constraints affect the algorithm.

Dimensional Constraints: Many objects have predetermined dimensions. Cars, road lanes, and chairs have a certain width. Stair steps and building floors have a certain height. Bowling balls and pool tables have a predetermined size. Without constraining the dimensions of the objects, the synthesis algorithm could easily generate roads too narrow to drive across, steps too tall to walk up, ceilings too close to the ground, and bowling balls too big to bowl. Dimensional constraints allow the user to fix the dimensions of the objects so that they are always sized realistically.

Algebraic Constraints: Some objects do not have predetermined dimensions, but instead must satisfy an algebraic relationship between their dimensions. An example might be that an object’s length must be twice its height. These constraints are especially useful for curved objects.

Incidence Constraints: Prior model synthesis techniques are limited to shapes which have only trihedral vertices. Trihedral vertices are vertices which are incident to three faces. As a result, there are many simple shapes such as a pyramid or an octahedron that previous model synthesis techniques cannot generate. To generate such shapes, we use additional incidence constraints.

Connectivity Constraint: Many objects look unnatural if they are not connected to a larger whole. One example is a road network. An loop of road looks unnatural if it is isolated from the all the other roads. All of the roads in a city are usually connected in some way. This defines a connectivity constraint which can be used to eliminate the possibility of isolated loops and create fully connected roads.

Large-Scale Constraints: The user might have a floor plan or a general idea of what the model should look

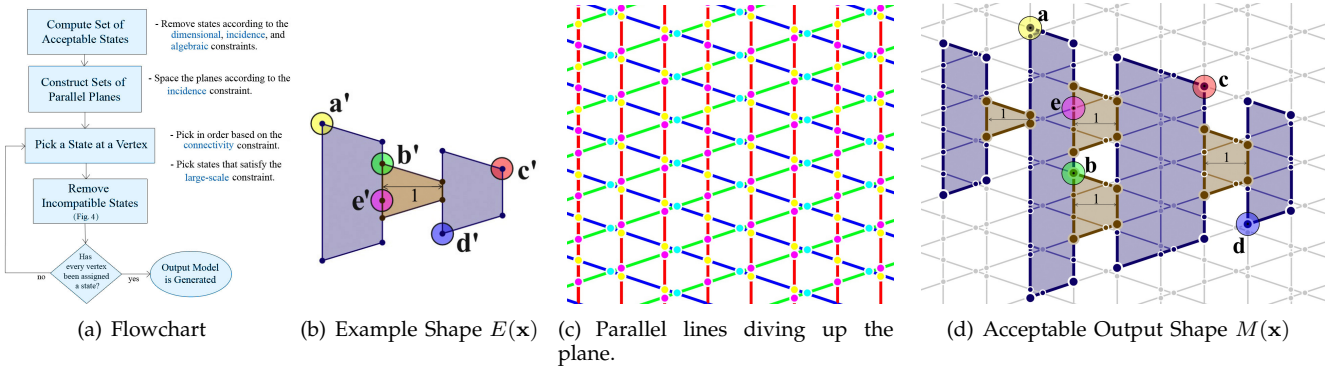


Fig. 2. (a) Overview flowchart showing how the constraints affect the algorithm. (b) From the input shape E , (c) sets of lines parallel to E intersect to form edges and vertices. (d) The output shape is formed within the parallel lines. For each selected point a , b , c , d , and e in M , there are points a' , b' , c' , d' , and e' in the example model E which have the same neighborhood. The models E and M contain two different kinds of object interiors shown in blue and brown. The brown object's width is constrained to be one line spacing width. The width of the blue object is not constrained. The objects are also constrained to be fully connected.

like on a macroscopic scale. For example, the user might want to build a city with buildings arranged in the shape of a circle or a triangle. The user can generate such a model by using large-scale constraints. These constraints are specified on a large volumetric grid where each voxel records which objects should appear within it.

4 CONSTRAINT-BASED APPROACH

In this section, we present our constraint-based synthesis algorithm.

4.1 Overview

We first discuss incidence constraints that specify that more than three faces are incident to a vertex in order to generate non-trihedral vertices. To add incidence constraints, we need a new way to describe the neighborhoods around non-trihedral vertices. We use Boolean expressions as explained in Section 4.2. These representations are used to determine which neighborhoods can be adjacent to one another. The vertices of the output are constructed where several planes intersect. Vertices incident to four faces require that four planes intersect which requires the planes to be spaced a particular way described in Section 4.4.

The Boolean expressions describing the states can also be used to apply dimensional constraints to the synthesized model. By disallowing any states that would permit the objects to stretch beyond its fixed dimensions, dimensional constraints are created as described in Section 4.5. Connectivity constraints are imposed in Section 4.6 by changing the order in which the states are assigned. Large-scale constraints are applied by changing the probabilities of the states that are assigned to each vertex, as described in Section 4.7. An algebraic constraint is described in Section 4.8.




Fig. 3. One edge and six vertex states are described using Boolean expressions of three half-spaces. In our algorithm, every neighborhood is represented by a Boolean expression.

4.2 Representing Neighborhoods with Boolean Expressions

In this section, we discuss how to describe states using Boolean expressions. The terms Boolean expression, neighborhood, and state can all be used interchangeably. The Boolean expressions are simply used to describe the neighborhoods. Our ultimate goal is to assign a neighborhood to each vertex of Figure 2(c). Each vertex has a set of possible neighborhoods that could be assigned to it and these are called states.

The incidence and the adjacency constraints are concerned with the neighborhoods surrounding points. To impose the adjacency constraint at non-trihedral vertices, we need a way to describe the neighborhoods there. Neighborhoods are represented using half-spaces which are related to the faces of the polyhedra. Every face has a plane that is parallel to it and that intersects the origin. This plane divides space into two half-spaces. The face's normal points into one half-space and away

from the other. Let us associate each face with the half-space that its normal points away from. These half-spaces can be used to describe every neighborhood of the polyhedra using a combination of Boolean operations. A few examples of these combinations are shown in Figure 3. For every point \mathbf{p} in E , there exists a Boolean expression that will produce a neighborhood identical to \mathbf{p} . This set of Boolean operations is found using the following method:

- 1) If \mathbf{p} is on a face and h_i is the half-space associated with the face, then h_i alone produces a neighborhood that is identical to \mathbf{p} . If \mathbf{p} is on a face whose normal points in the opposite direction, then h_i^C describes the neighborhood around \mathbf{p} .
- 2) If \mathbf{p} is on an edge whose two adjacent faces are associated with the two half-spaces h_i and h_j , then the neighborhood around \mathbf{p} is described by $h_i \cup h_j$ if the edge has a reflex angle and $h_i \cap h_j$ if it does not.
- 3) If \mathbf{p} is on a vertex, then the procedure for computing its neighborhood's Boolean expression is more complex. Every face that intersects \mathbf{p} is on a plane. Let us take all the faces that intersect \mathbf{p} and use all of their planes to divide the space into cells. An example of this is shown in Figure 5. Each cell is the intersection of several half-spaces. Since the planes all intersect \mathbf{p} , every cell has points in the neighborhood of \mathbf{p} . For each cell, we determine if the points within the cell and within the neighborhood of \mathbf{p} are inside or outside the polyhedron E . We take the union of all cells which have points inside the polyhedron and this is the Boolean expression that represents the neighborhood surrounding \mathbf{p} . Each cell is the intersection of several half-spaces and so the neighborhood at \mathbf{p} is represented as a union of intersections. These expressions can often be simplified using familiar rules of Boolean algebra such as $(h_i \cap h_j) \cup (h_i \cap h_j^C) = h_i$. Simplified Boolean expressions for various states are shown in Figures 3 and 5.

This method gives us a Boolean expression describing how a polyhedra intersects a neighborhood at any point \mathbf{p} . However, more than one polyhedra might intersect at the same point. For example, see the points \mathbf{b} and \mathbf{e} of Figure 2(d). When multiple objects intersect, we can compute a Boolean expression for each object and then combine all the expressions into one. Let k_1 and k_2 be two object interiors and let b_1 and b_2 be two Boolean expressions. The notation $k_1 \cdot b_1 + k_2 \cdot b_2$ will be used to describe a neighborhood which contains the object k_1 at the points b_1 and object k_2 at b_2 . For example, $1 \cdot h_1 \cap h_2^C + 3 \cdot h_2$ describes a neighborhood where an edge $h_1 \cap h_2^C$ of object 1 touches a face of h_2 of object 3.

4.3 Evaluating Boolean Expressions along Edges

In the previous section, we discuss how to describe every neighborhood or every state as a Boolean expression, but

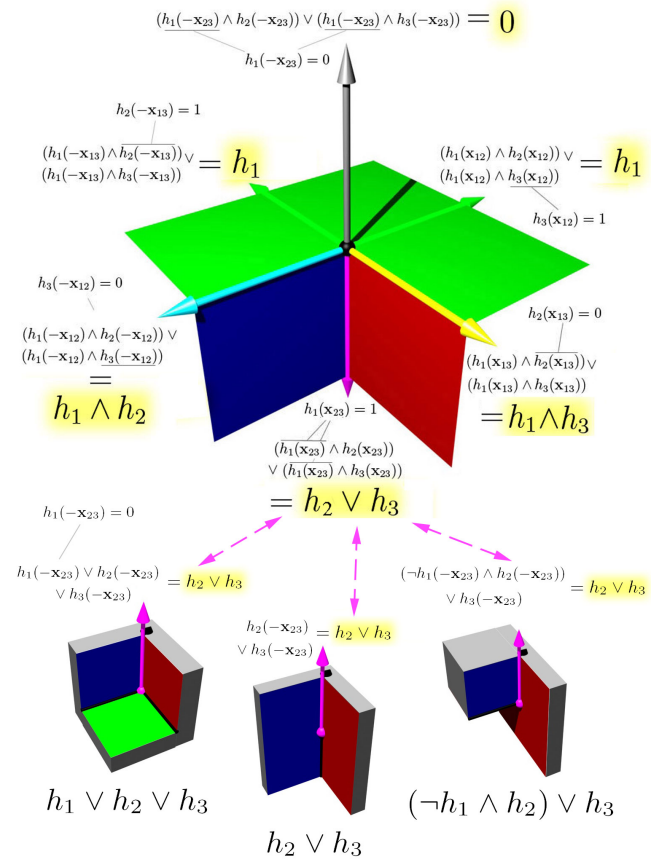


Fig. 4. We can evaluate the states to figure out which states can be adjacent to them in various directions. This shows the state $(h_1 \wedge h_2) \vee (h_1 \wedge h_3)$ evaluated in six directions. We assume that \mathbf{x}_{23} points down and is inside the h_1 half-space so that $h_1(\mathbf{x}_{23}) = 1$ and $h_1(-\mathbf{x}_{23}) = 0$. In the \mathbf{x}_{23} direction, the state $(h_1 \wedge h_2) \vee (h_1 \wedge h_3)$ evaluates to $h_2 \vee h_3$ which is shown as the magenta edge. A state can only be underneath the state $(h_1 \wedge h_2) \vee (h_1 \wedge h_3)$ if it shares the same $h_2 \vee h_3$ edge. We test every state to see which ones evaluate to $h_2 \vee h_3$ in the $-\mathbf{x}_{23}$ direction. Three acceptable states are shown: $h_1 \vee h_2 \vee h_3$, $h_2 \vee h_3$, and $(\neg h_1 \wedge h_2) \vee h_3$.

we still need to determine which states can be next to one another. The Boolean expressions can be thought about in more than one way. We have been thinking in terms of union and intersection of sets, but we could replace them with OR and AND operations of functions. Each half-space has a characteristic function $h_i(\mathbf{x})$ which evaluates to 1 if \mathbf{x} is inside the half-space and to 0 if \mathbf{x} is in the opposite half-space. However, a third possibility is that \mathbf{x} intersects the plane dividing the two half-spaces. In this case, we do not evaluate $h_i(\mathbf{x})$ as 0 or 1, but leave it as the symbol h_i . This symbolic representation provides a convenient way to determine how the states connect together.

Since the evaluation depends on which planes a point \mathbf{x} intersects, we keep track of the planes that a point




Fig. 5. Even complex vertices can be described using a Boolean expression of half-spaces. We can take all the faces that intersect the vertex and use all of the planes the faces are on to divide up the space into regions that are the intersection of several half-spaces. The regions labeled in the figure are inside the polyhedron. The vertex state is described as the union of all these interior regions. This vertex is simplified to $(h_3 \wedge (h_1 \vee h_2)) \vee (h_4 \wedge (h_1 \vee h_5))$.

intersects using subscripts. According to this notation, the point $\mathbf{x}_{12} = \mathbf{n}_1 \times \mathbf{n}_2$ is on the planes of h_1 and h_2 .

The Boolean expressions can describe many different neighborhoods or states, including vertices, edges, and faces. When we evaluate the expression at a point \mathbf{x} , essentially we get a new state that describes what we would encounter if we travel away from a neighborhood in the direction of \mathbf{x} . If the expression evaluates to 0, we travel into empty space. If it evaluates to 1, we travel into an object's interior. If it evaluates to h_1 , we travel onto a face. If it evaluates to something like $h_1 \wedge h_2$ or $h_1 \vee h_2$, then we travel onto an edge. We determine if two states can be adjacent to one another by evaluating the states in opposite directions and checking if their evaluations are identical. In Figure 4, the state $(h_1 \wedge h_2) \vee (h_1 \wedge h_3)$ evaluates in the $-z$ direction to $h_2 \vee h_3$. Any state that evaluates to $h_2 \vee h_3$ in the $+z$ direction can be beneath the state $(h_1 \wedge h_2) \vee (h_1 \wedge h_3)$. Three examples of such states are shown in Figure 4.

The Boolean expressions may contain more than one object interior. In this case, we evaluate each object interior separately and combine the results. For example, the expression $1 \cdot (h_1 \wedge \neg h_2) + 3 \cdot h_2$ is used to describe a neighborhood in which an edge $h_1 \wedge \neg h_2$ of object 1 touches a face h_2 of object 3. If we evaluate the expression at the point \mathbf{x}_{23} and if $h_1(\mathbf{x}_{23}) = 1$, then we would compute $1 \cdot h_1(\mathbf{x}_{23}) \wedge \neg h_2(\mathbf{x}_{23}) + 3 \cdot h_2(\mathbf{x}_{23}) = 1 \cdot \neg h_2 + 3 \cdot h_2$. This means that if we travel in the direction \mathbf{x}_{23} we will encounter two faces up against each other. One face is from object 1 and the other is from object 3.

4.4 Spacing the Planes

For each vertex of Figure 2(c), we compute a list of possible states. Each state corresponds to a neighborhood which can be described by a Boolean expression. We first find all states found in the input model. Each face, each edge, each vertex, and each object interior is a different

state. However, only a small fraction of these states are acceptable at each vertex of Figure 2(c). To be acceptable, every half-space used in the state's Boolean expression must be associated with a plane that the vertex intersects. A trihedral vertex is described using three half-spaces; three half-spaces require three planes to intersect; and those three planes will intersect somewhere. But non-trihedral vertices pose a difficult problem since they can only appear where four or more planes intersect and four planes may never intersect. For these cases, we need to choose the plane spacing so that four planes intersect.

We first discuss how to compute the points where three planes intersect and then discuss how to get a fourth plane to intersect the same points. Let $\mathbf{n}_1, \mathbf{n}_2$, and \mathbf{n}_3 be the normals of three sets of planes. Within each set of planes, the planes are parallel and evenly spaced. Let s_1, s_2 , and s_3 be the spacing between the planes within each set. The first and second sets of planes intersect along lines that point in the $\mathbf{n}_1 \times \mathbf{n}_2$ direction. If \mathbf{p} is a point on one of the lines, then \mathbf{p}' is on the same line if $\mathbf{p}' - \mathbf{p} = (\mathbf{n}_1 \times \mathbf{n}_2)t$ for some scalar t . If \mathbf{p} also intersects a plane from the third set, then \mathbf{p}' intersect a plane if $\mathbf{n}_3 \cdot (\mathbf{p}' - \mathbf{p}) = c_3 s_3$ for some integer $c_3 \in \mathbb{Z}$. Solving for t , we find that $t = \frac{c_3 s_3}{\mathbf{n}_3 \cdot (\mathbf{n}_1 \times \mathbf{n}_2)}$ and, therefore,

$$\mathbf{p}' = \mathbf{p} + c_3 s_3 \frac{\mathbf{n}_1 \times \mathbf{n}_2}{\mathbf{n}_3 \cdot (\mathbf{n}_1 \times \mathbf{n}_2)} \quad (2)$$

for some $c_3 \in \mathbb{Z}$. This gives us a set of points along the $\mathbf{n}_1 \times \mathbf{n}_2$ direction where the three planes intersect. The same argument can also be applied to the $\mathbf{n}_1 \times \mathbf{n}_3$ and $\mathbf{n}_2 \times \mathbf{n}_3$ directions. Three planes intersect at the points

$$\mathbf{p}' = \mathbf{p} + \frac{c_1 s_1 \mathbf{n}_2 \times \mathbf{n}_3}{\mathbf{n}_1 \cdot (\mathbf{n}_2 \times \mathbf{n}_3)} + \frac{c_2 s_2 \mathbf{n}_1 \times \mathbf{n}_3}{\mathbf{n}_2 \cdot (\mathbf{n}_1 \times \mathbf{n}_3)} + \frac{c_3 s_3 \mathbf{n}_1 \times \mathbf{n}_2}{\mathbf{n}_3 \cdot (\mathbf{n}_1 \times \mathbf{n}_2)} \quad (3)$$

for any $c_1, c_2, c_3 \in \mathbb{Z}$. Each different combination of c_1, c_2 , and c_3 gives us a different intersection point and the resulting intersection points form a 3D lattice. The three planes always intersect regardless of how they are spaced, but it is much more difficult to get four planes to intersect. If \mathbf{p} intersects a plane from the fourth set, then \mathbf{p}' also does if $\mathbf{n}_4 \cdot (\mathbf{p}' - \mathbf{p}) = c_4 s_4$ for some integer $c_4 \in \mathbb{Z}$. The point \mathbf{p}' intersects the fourth set of planes if

$$s_4 = s_1 \frac{c_1 \mathbf{n}_4 \cdot (\mathbf{n}_2 \times \mathbf{n}_3)}{c_4 \mathbf{n}_1 \cdot (\mathbf{n}_2 \times \mathbf{n}_3)} + s_2 \frac{c_2 \mathbf{n}_4 \cdot (\mathbf{n}_1 \times \mathbf{n}_3)}{c_4 \mathbf{n}_2 \cdot (\mathbf{n}_1 \times \mathbf{n}_3)} + s_3 \frac{c_3 \mathbf{n}_4 \cdot (\mathbf{n}_1 \times \mathbf{n}_2)}{c_4 \mathbf{n}_3 \cdot (\mathbf{n}_1 \times \mathbf{n}_2)} \quad (4)$$

for some $c_4 \in \mathbb{Z}$. Equation (4) describes an algebraic relationship that must be satisfied for four planes to intersect. This relationship represents multiple equations that need to be solved since each combination of c_1, c_2 , and c_3 produces another equation. If we solve this equation for the combinations $(c_1, c_2, c_3, c_4) = (1, 0, 0, 1)$ and $(0, 1, 0, 1)$ and $(0, 0, 1, 1)$, then it will hold for any combination of c_1, c_2 , and c_3 . Essentially, we have three equations and four unknowns s_1, s_2, s_3 , and s_4 . By solving for these

linear equations, we produce a 3D lattice of points where a non-trihedral vertex state may appear. However, this only takes care of a single non-trihedral vertex state. There may be more of these states in the input and they would each require more equations to be solved. There are even more difficult vertex states to handle like the vertex shown in Figure 5 which involve five half-spaces. These require solving more linear equations.

In the end, we may have an underconstrained or an overconstrained set of linear equations. An overconstrained set of equations occurs when the input model does not fit well within a lattice. One example of an input shape that produces overconstrained equations is a five-sided pyramid. These overconstrained equations can be handled in several ways. One approach is to add many more planes, but this increases the computational cost of the overall algorithm. Another approach might be to modify the normals just enough that the shapes better fit on a lattice, but not so much that the normals significantly change the results. A third option is to leave a few of the equations unsatisfied. When this happens, non-trihedral vertices will be generated at fewer locations, but this might be adequate to produce a good final result.

4.5 Dimensional Constraints

We would like to give the user greater control over the dimensions of the output. The user should be able to control if an object can scale in a particular direction. For example, a user might specify that a road must have a particular width. Along its width, the road cannot scale, but along its length, the road can scale to any length. The ability to fix the dimensions of some objects is important for creating realistic models.

Since the objects are created on sets of evenly spaced planes, the lengths of each object must be an integer multiple of the plane spacing. Objects with non-integer dimensions like for example 1.5-plane spaces can pose a problem. To deal with these objects the planes could be spaced more closely. If they are spaced twice as close, an object that was 1.5-plane spaces wide would become three planes wide which is a round number. Often there is an even simpler solution since objects with dimensional constraints are often next to objects without them and the two objects can be attached together to produce a round number. For example, it might be possible to combine an object 1.5 spaces wide with 0.5 spaces of empty space to produce an object two plane spaces wide which is a round number.

Even though objects may be two, three or more plane-spaces wide, we only need to consider the issue of how to force an object to be exactly one-space wide since we can easily create objects exactly two or three spaces wide simply by attaching a few one-space wide objects together.

Figure 6 shows a simple example of how this constraint is imposed. The objects can never grow wider




Fig. 6. Dimensional Constraint. To create objects that are only one plane space wide horizontally, we disallow any states which pass through the vertical h_2 planes such as h_1 or $\sim h_1$.

than one plane-space if every time they intersect a plane they stop. To stop their growth, we disallow all vertex states in which the object passes through the plane. The object is on both sides of the plane if $h_2 \wedge b \neq 0$ and $\sim h_2 \wedge b \neq 0$ where h_2 is the half-space parallel to the plane and b is the Boolean expression describing the vertex state. By removing all states where $h_2 \wedge b \neq 0$ and $\sim h_2 \wedge b \neq 0$, we guarantee that the objects do not grow more than one plane-space wide.

4.6 Connectivity Constraints

In many applications, controlling the connections between objects is important. For example, this is important when creating urban models with roads. In most cities, one could choose any two points on a road map and find a path that connects them. However, model synthesis algorithms could generate isolated loops or cycles of road networks that are not connect to each other. This problem can be addressed by changing the order in which the states are assigned. We begin by choosing a starting location at random and creating an object (e.g. road) there. Then the roads are all grown out from this initial seed. This means that we only assign road states to vertices that are already next to a road. By growing out from a single seed, the generated roads are fully connected.

A fully connected object is just one of several options to consider. One alternative is to not use seeds at all and to assign the states in any order. This is useful when the user wants to create many isolated objects. A third option fits in between the other two. The user might not want everything to be connected, but might not want many small isolated objects either. The user may want a few large isolated objects. To accomplish this, everything could be grown out not from a single seed, but from multiple seeds.

4.7 Large-Scale Constraints

We would also like to give the user more control over the large-scale structure of the output. The user might have a general idea of where certain types of objects should appear. Each object has a particular probability that it will appear at any location in space. Generally, we choose to give each state an equal probability of being chosen, but we could easily modify the probabilities so that they are higher for any particular objects the user wants to appear within some areas. The user could even set some probabilities to be zero in some places. If a state's probability drops to zero, we can remove it entirely and then propagate the removal as usually done when assigning states (see Section 3.2). By changing these probabilities, we can create cities and other structures in the shape of various symbols and other objects. We can also generate multiple outputs, evaluate how well they match the user's desired goal, and select the best output.

4.8 Algebraic Constraints and Bounding Volumes

The model synthesis algorithm creates a set of parallel planes for every distinct normal of the input. As a result, handling curved input models with many distinct normals are computationally expensive because of the large number of planes that would have to be created. However, the number of distinct normals can be greatly reduced by using bounding boxes and other bounding volumes in place of complex objects. The algorithm could be run using the bounding volumes in place of the input model and complex objects can be substituted back into the output model M after it is generated.

There are several alternative ways the user can constrain the dimensions. The object's dimensions could scale freely in a direction or be fixed (see Section 4.5). A third option is to let an object scale, but to require that it must scale uniformly in two or three directions. For example, the cylinder in Figure 7 only remains cylindrical if its x and y coordinates scale uniformly $s_x = s_y$. It is free to stretch along the z -coordinate by any amount. To get its x and y coordinates to scale equally, we can place a bounding box around the cylinder and the cut the box into two halves along the diagonal creating two triangular prisms shown in Figure 7. Since model synthesis scales triangular objects uniformly in two dimensions, the output will be scaled identically in x and y , $s_x = s_y$ and the cylinder can be substituted back in the shape.

The user may want to be even more restrictive and require the scalings be uniform in all directions $s_x = s_y = s_z$. For example, the dome in Figure 7 remains spherical only in this case. This can be accomplished by placing a bounding box around the sphere and cutting off a tetrahedron as shown in Figure 7. Since model synthesis scales tetrahedra uniformly in all directions, the output will create a uniformly scaled copy of the bounding box.

	Input Size (polygons)	Output Size (polygons)	Time (minutes)
Oil Platform	60	1,377	0.5
Domes	21	324	0.1
Buildings	116	2,230	1.4
Spaceships	168	4,164	0.6
Roads	126	6,888	0.2
Roller Coaster	124	1,376	1.8
GPM	365	7,527	3.5

TABLE 1

Complexity of the input and output models and computation time for various results computed on a 2.8 GHz single-core PC.

5 RESULTS

Figures 1, 7 - 12 show a variety of models that were generated using our algorithm. The generated models are large and detailed and it would be quite difficult to model them manually using a CAD or authoring system. The models each satisfy multiple constraints which depend on the application. Dimensional constraints are used in Figure 1 to give the platforms and beams a fixed thickness. They are also used to constrain the width of the road in Figure 10, the width of the spacecrafts in Figure 9, and the width of the roller coaster track in Figure 11. Incidence constraints are repeatedly used in Figure 8 to create architecture with four faces touching at a single vertex. Connectivity constraints are used in Figures 1, 8, and 9 to grow the objects out from a few seeds and this controls the distribution of the objects so they are not all crowded together. The roads in Figure 10 are fully connected to a single seed. In Figure 9, the parts of the spaceships are connected by beams and have gaps in between. Figure 9 demonstrates that model synthesis can generate shapes which have a high genus. Bounding volumes were used in Figures 1 and 9 to generate curved objects. Algebraic constraints were used in Figure 7. A large-scale constraint is used in Figure 12 to generate several different types of objects in the form of the characters "GPM".

Each of the models was generated without requiring much effort from the user. The input models E are composed of only a few dozen polygons. Each of the constraints can be specified by only changing a few parameters or in the case of Figure 12 an image of the letters "GPM" Table 1 shows the computation time for modeling each shape and the size of the input and output models. The size is given in terms of the polygon count of the bounding volumes. All of the displayed images include artistic decorations to the vertices and edges and some include complex objects that were generated from bounding volumes. The polygon count does not include any of these decorations. The computation time depends both on the output size and on which input model is used. Some models can be computed much more quickly than others. The road model (Figure 10) can be computed quickly because it is flat and does not

really use all three spatial dimensions. The ‘GPM’ model takes the longest time to compute because it uses several different input models including a spaceship model and several building models.

6 LIMITATIONS

The amount of time and memory that model synthesis needs depends on the number of vertices. Vertices are generated wherever three or more planes intersect. The number of planes depends on the number of distinct face normals. If there are n distinct normals and m parallel planes for each normal, there could be up to $O(n^3m^3)$ vertices. The number of distinct normals can be reduced by using bounding volumes, but only to a certain extent since bounding volumes can oversimplify some objects. This makes generating curved objects using model synthesis especially difficult [30]. A related problem is that it is difficult to generate both large and small objects simultaneously. Small objects require closely spaced planes while large objects require large volumes which together means that many planes must be created.

Like most procedural modeling techniques, model synthesis is designed to work on objects that are self-similar. Model synthesis works best on objects with parts that identically match. Objects without identical parts can be used with model synthesis, but they often produce results that match the input too closely. Since model synthesis is meant for digital entertainment and gaming, we assume that objects in the input are free from significant errors in the vertex positions. Model synthesis works efficiently on man-made structures that can be represented with a few planar faces, but it has difficulty with organic and curved shapes.

Another limitation is that the objects often need to have a grid structure. The grid is a necessary part of some of the constraints. The dimensional constraint assumes the dimensions fit on a grid. The incidence constraint assumes that the vertices can be fit onto a grid. The structure of the grid depends on the plane spacing which can be altered to accommodate some shapes, but not all shapes as explained in Section 4.4. Some shapes may produce an overconstrained set of equations when using Equation 4. Several strategies for dealing with this problem were discussed in Section 4.4, but each of them has downsides.

7 COMPARISON

7.1 Prior Model Synthesis Algorithms

Previous model synthesis techniques [28] only use the adjacency constraint. Using prior algorithms without the dimensional and algebraic constraints, most of the results would appear distorted and unnatural. Without the connectivity constraint, the resulting models would contain mostly small crowded objects. Without the incidence constraints, none of the buildings in Figure 8 would

be generated. The large-scale constraint is required to generate the pattern in Figure 12.

7.2 Other Procedural Modeling Techniques

Each of the model synthesis algorithms are semi-automatic. The user must perform several tasks. The difficulty of these tasks depends on the type of object that is being modeled. Objects that fit on a grid or contain mostly flat surfaces are relatively easy to generate using model synthesis. These types of objects are often man-made objects and are frequently found in the architectural domain. But other shapes are more difficult to generate using model synthesis including many natural and organic shapes. Organic shapes are difficult to generate with model synthesis because they do not fit on a grid and have many distinct normals. While model synthesis is not useful for generating every type of objects, it offers benefits over other procedural methods for many classes of objects. The user has a relatively simple and straightforward objective: to find or to create an example model and describe a set of constraints on the output.

In contrast, the user’s objective is less simple and straightforward for many existing procedural modeling techniques. Many techniques require the user to construct a grammar. Given the shape of an object the user wants to model, there may not be a straightforward procedure for constructing the rules of a grammar that could generate a similar shape. Grammars are constructed through some human ingenuity and through trial and error. The grammars themselves can be complicated, even when they describe simple shapes.

The model synthesis algorithms are easier to understand from a user’s perspective. The user does not need to know anything about grammars or the inner mechanics of the algorithm itself. The user only needs to know a few basic facts about the algorithm. The user deals only with the input and output models and the constraints. The user needs to avoid creating curved surfaces or to put bounding volumes around them. Once a suitable example model has been created it is easy to modify it as needed. Its parts can easily be rearranged using standard 3D modeling programs.

Most procedural modeling techniques are aimed at modeling specific classes of objects such as urban buildings [31], truss structures [39], fractals, and landscapes [32]. But many interesting structures lie outside of these classes of objects including spaceships, castles, oil platforms, plumbing, and roller coasters, just to name a few. Since model synthesis is a more general technique, it is especially useful for modeling objects that cannot be generated easily using other techniques.

A close connection between model synthesis and context-sensitive grammars is demonstrated in Section 7.3. Both methods can be used to accomplish the same goal. The set of acceptable models can be represented by a grammar, but it is different in several ways from grammars that are typically used in procedural modeling. The

location of empty space is recorded in model synthesis. Empty space is not explicitly recorded in most grammar-based techniques. It is determined by checking that all of the objects are absent. Model synthesis algorithms are also good at avoiding self-intersections which is part of the adjacency constraint. The grammars found in other techniques may need to be carefully constructed so that self-intersections do not occur. Another task that model synthesis is particularly good at is in creating closed paths such as Figure 10.

A method [5] closely related to model synthesis automatically detects symmetrical parts and derives a shape grammar. This method complements work in model synthesis since the user would not need to label symmetrical parts if this were computed automatically. Like other shape grammars, the derived shape grammar can not always anticipate and avoid self-intersections and may have difficulty forming closed paths.

Model synthesis algorithms are good at creating geometric detail at a particular scale, but not at multiple scales. For example, it is difficult for model synthesis to create geometric detail at the scale of a building and at the scale of the building's window or door knob simultaneously. Other grammar-based methods [31], [25], [32] create geometric detail at multiple scales more easily.

7.3 Relationship between model synthesis and Context Sensitive Grammars

The problem of deciding if a string is part of the context sensitive language L can be reduced to a problem of deciding if a model containing that string satisfies the adjacency constraint. A model is *consistent* if it satisfies the adjacency constraint.

For every context-sensitive language L , there is a linear-bounded automaton that accepts L . It can be shown that all of the actions of a linear-bounded automaton A accepting a string can be described within a consistent model.

For any linear bounded automaton, an adjacency constraint can be constructed such that the model that is generated will reproduce the actions of the automaton. Each row of the model records the symbols on the tape, the location of the tape head, and the state of A . For example, suppose the problem is to determine if the string 'aabbcc' is in the language $\{a^i b^j c^k \mid i \geq 1\}$. The tape would initially contain this input string along with two symbols '<' and '>' to mark the start and end of the tape. The width of the model M is equal to the width of the tape. The first row of the model would contain the labels

Row 1: < (a,q₀) a b b c c >

The label (a,q₀) is used to indicate that the automaton A is in its initial state q_0 and the tape head is reading the 'a' symbol. For every state q and every symbol in the tape alphabet s , there is a label (q, s) . Suppose that when A is in state q_0 and is reading symbol 'a' that it responds by

printing the symbol 'd' onto the tape, switching to state q_1 , and remaining stationary. Then the model's next row would be

Row 2: < (d,q₁) a b b c c >

An adjacency constraint can be constructed which would guarantee that this row would appear beneath row 1. The constraint is constructed to allow only one possible option at every location which is exactly the option that the automaton would choose. The constraint can be constructed so that the tape head also moves to the left or to the right. The problem of determining if a string is part of a context-sensitive language can be decided by determining if such a model can be completed acceptably. This gives an overview of the proof. More details can be found in [27].

We have discussed how a context-sensitive grammar problem can be reduced to a model synthesis problem. The reverse is also true. A context-sensitive grammar can be used to generate consistent models. For every example model E there is a linear-bounded automaton that can examine all of the labels in the model and verify that all adjacency labels satisfy the adjacency constraint. This topic is explored in more detail in [27].

8 CONCLUSION AND FUTURE WORK

We have presented several major improvements to model synthesis that allow the user to more effectively control the output. We enable the user to fix dimensions of objects, to specify a large-scale structure of the output, to produce connected results, to add bounding volumes, to have multiple object interiors, and to generate shapes with complex vertex states. Further work is needed to improve the efficiency of model synthesis, especially when generating large and small objects together. More work is needed for handling curved objects beyond using bounding volumes [30]. One important constraint that is still missing is one to create symmetrical objects.

ACKNOWLEDGEMENTS

We would like to thank the reviewers for their comments. This work was supported in part by ARO Contract W911NF-04-1-0088, NSF awards 0636208, 0917040 and 0904990, and DARPA/RDECOM Contract WR91CRB-08-C-0137.

BIOGRAPHY

Paul Merrell is currently a Postdoctoral Scholar at Stanford University. He received his Ph.D. in Computer Science at the University of North Carolina at Chapel Hill in 2009. He also received a Bachelor's and Master's degree in Electrical Engineering at Brigham Young University. Paul Merrell has worked on many projects in computer graphics and computer vision and spent several summers working at EA Games and Raytheon Missile Systems.

Dinesh Manocha is currently the Phi Delta Theta/Mason Distinguished Professor of Computer Science at the University of North Carolina at Chapel Hill. He received his Ph.D. in Computer Science at the University of California at Berkeley 1992. He has received Junior Faculty Award, Alfred P. Sloan Fellowship, NSF Career Award, Office of Naval Research Young Investigator Award, Honda Research Initiation Award, Hettleman Prize for Scholarly Achievement. Along with his students, Manocha has also received 12 best paper & panel awards at the leading conferences on graphics, geometric modeling, visualization, multimedia and high-performance computing. He is an ACM Fellow.

Manocha has published more than 280 papers in the leading conferences and journals on computer graphics, geometric computing, robotics, and scientific computing. He has also served as a program committee member and program chair for more than 75 conferences in these areas, and editorial boards of many leading journals. Some of the software systems related to collision detection, GPU-based algorithms and geometric computing developed by his group have been downloaded by more than 100,000 users and are widely used in the industry. He has supervised 18 Ph.D. dissertations.

REFERENCES

- [1] H. Ault, "Using geometric constraints to capture design intent," *Journal for Geometry and Graphics*, vol. 3, no. 1, pp. 39–47, 1999.
- [2] D. G. Aliaga, C. A. Vanegas, and B. Beneš, "Interactive example-based urban layout synthesis," *ACM Trans. Graph.*, vol. 27, no. 5, pp. 1–10, 2008.
- [3] C. Barnes, E. Shechtman, A. Finkelstein, and D. B. Goldman, "PatchMatch: A randomized correspondence algorithm for structural image editing," *ACM Transactions on Graphics (Proc. SIGGRAPH)*, vol. 28, no. 3, Aug. 2009.
- [4] P. Bhat, S. Ingram, and G. Turk, "Geometric texture synthesis by example," in *SGP '04: Symposium on Geometry processing*. New York, NY, USA: ACM, 2004, pp. 41–44.
- [5] M. Bokeloh, M. Wand, H. Seidel, "A Connection between Partial Symmetry and Inverse Procedural Modeling," in *Proc. of ACM SIGGRAPH*, 2010.
- [6] W. Bouma, I. Fudos, C. Hoffmann, J. Cai, and R. Paige, "A geometric constraint solver," *Computer-Aided Design*, vol. 27, no. 6, pp. 487–501, 1995.
- [7] G. Chen, G. Esch, P. Wonka, P. Müller, and E. Zhang, "Interactive procedural street modeling," *ACM Trans. Graph.*, vol. 27, no. 3, 2008.
- [8] X. Chen, B. Neubert, Y.-Q. Xu, O. Deussen, and S. B. Kang, "Sketch-based tree modeling using markov random field," *ACM Trans. Graph.*, vol. 27, no. 5, pp. 1–9, 2008.
- [9] B. Cutler, J. Dorsey, L. McMillan, M. Müller, and R. Jagnow, "A procedural approach to authoring solid models," *ACM Trans. Graph.*, vol. 21, no. 3, pp. 302–311, 2002.
- [10] A. A. Efros and W. T. Freeman, "Image quilting for texture synthesis and transfer," in *SIGGRAPH '01*. New York, NY, USA: ACM, 2001, pp. 341–346.
- [11] A. A. Efros and T. K. Leung, "Texture synthesis by non-parametric sampling," in *IEEE International Conference on Computer Vision*, Corfu, Greece, September 1999, pp. 1033–1038.
- [12] U. Flemming, "More than the sum of parts: the grammar of queen anne houses," *Environment and Planning B: Planning and Design*, vol. 14, no. 3, pp. 323–350, May 1987.
- [13] A. Fournier, D. Fussell, and L. Carpenter, "Computer rendering of stochastic models," *Commun. ACM*, vol. 25, no. 6, pp. 371–384, 1982.
- [14] T. Funkhouser, M. Kazhdan, P. Shilane, P. Min, W. Kiefer, A. Tal, S. Rusinkiewicz, and D. Dobkin, "Modeling by example," *SIGGRAPH '04*, 2004.
- [15] C. M. Hoffmann and J. R. Rossignac, "A road map to solid modeling," *IEEE Transactions on Visualization and Computer Graphics*, vol. 2, no. 1, pp. 3–10, 1996.
- [16] C. M. Hoffmann, A. Lomonosov, and M. Sitharam, "Geometric constraint decomposition," in *Geometric Constraint Solving*, B. Bruderlin and D. Roller, Eds. Springer-Verlag, 1998, pp. 170–195.
- [17] T. Igarashi, S. Matsuoka, and H. Tanaka, "Teddy: a sketching interface for 3d freeform design," in *Proc. of ACM SIGGRAPH '99*. New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 1999, pp. 409–416.
- [18] G. A. Kramer and B. B. Qh, "Solving geometric constraint systems." MIT Press, 1992, pp. 708–714.
- [19] J. Kopf, C.-W. Fu, D. Cohen-Or, O. Deussen, D. Lischinski, and T.-T. Wong, "Solid texture synthesis from 2d exemplars," *ACM Trans. Graph.*, vol. 26, no. 3, p. 2, 2007.
- [20] V. Kwatra, A. Schödl, I. Essa, G. Turk, and A. Bobick, "Graphcut textures: image and video synthesis using graph cuts," in *SIGGRAPH '03*. New York, NY, USA: ACM, 2003, pp. 277–286.
- [21] J. Legakis, J. Dorsey, and S. Gortler, "Feature-based cellular texturing for architectural models," in *Proc. Of ACM SIGGRAPH '01*, 2001, pp. 309–316.
- [22] A. Lindenmayer, "Mathematical models for cellular interactions in development i. filaments with one-sided inputs," *Journal of Theoretical Biology*, vol. 18, no. 3, pp. 280–299, March 1968.
- [23] B. B. Mandelbrot, *The Fractal Geometry of Nature*. W. H. Freeman, August 1982.
- [24] N. J. Mitra, L. Guibas, and M. Pauly, "Partial and approximate symmetry detection for 3d geometry," in *ACM Transactions on Graphics*, vol. 25, no. 3, 2006, pp. 560–568.
- [25] R. Měch and P. Prusinkiewicz, "Visual models of plants interacting with their environment," in *Proc. Of ACM SIGGRAPH '96*, 1996, pp. 397–410.
- [26] P. Merrell, "Example-based model synthesis," in *IBD '07: Symposium on Interactive 3D graphics and games*. ACM Press, 2007, pp. 105–112.
- [27] P. Merrell, "Model synthesis," Ph.D. dissertation, University of North Carolina at Chapel Hill, 2009.
- [28] P. Merrell and D. Manocha, "Continuous model synthesis," *Proc. of ACM SIGGRAPH ASIA '08*, 2008.
- [29] P. Merrell and D. Manocha, "Constraint-based model synthesis," in *SPM '09: 2009 SIAM/ACM Joint Conference on Geometric and Physical Modeling*. New York, NY, USA: ACM, 2009, pp. 101–111.
- [30] P. Merrell and D. Manocha, "Example-based curve synthesis," in *Computers & Graphics*. (To Appear).
- [31] P. Müller, P. Wonka, S. Haegler, A. Ulmer, and L. V. Gool, "Procedural modeling of buildings," *ACM Trans. Graph.*, vol. 25, no. 3, pp. 614–623, 2006.
- [32] F. K. Musgrave, C. E. Kolb, and R. S. Mace, "The synthesis and rendering of eroded fractal terrains," in *Proc. Of ACM SIGGRAPH '89*, 1989, pp. 41–50.
- [33] A. Nealen, T. Igarashi, O. Sorkine, and M. Alexa, "Fibermesh: designing freeform surfaces with 3d curves," *Proc. of ACM SIGGRAPH '07*, vol. 26, no. 3, p. 41, 2007.
- [34] Y. Parish and P. Müller, "Procedural modeling of cities," *Proc. of ACM SIGGRAPH '01*, pp. 301–308, New York, NY, USA, 2001.
- [35] H. Pottmann, Y. Liu, J. Wallner, A. Bobenko, and W. Wang, "Geometry of multi-layer freeform structures for architecture," *Proc. Of ACM SIGGRAPH '07*, 2007.
- [36] P. Prusinkiewicz, A. Lindenmayer, and J. Hanan, "Development models of herbaceous plants for computer imagery purposes," *SIGGRAPH Comput. Graph.*, vol. 22, no. 4, pp. 141–150, 1988.
- [37] P. Prusinkiewicz, L. Mündermann, R. Karwowski, and B. Lane, "The use of positional information in the modeling of plants," in *Proc. Of ACM SIGGRAPH '01*, 2001, pp. 289–300.
- [38] L. Quan, P. Tan, G. Zeng, L. Yuan, J. Wang, and S. B. Kang, "Image-based plant modeling," *ACM Trans. Graph.*, vol. 25, no. 3, pp. 599–604, 2006.
- [39] J. Smith, J. Hodgins, I. Oppenheim, and A. Witkin, "Creating models of truss structures with optimization," *ACM Trans. Graph.*, vol. 21, no. 3, pp. 295–301, 2002.
- [40] G. Stiny and W. J. Mitchell, "The palladian grammar," *Environment and Planning B: Planning and Design*, vol. 5, no. 1, pp. 5–18, January 1978.
- [41] P. Torrens, David, and Sullivan, "Cellular automata and urban simulation: where do we go from here?" *Environment and Planning B: Planning and Design*, vol. 28, no. 2, pp. 163–168, March 2001.
- [42] G. Turk, "Texture synthesis on surfaces," in *SIGGRAPH '01*. New York, NY, USA: ACM, 2001, pp. 347–354.
- [43] C. Vanegas, D. Aliaga, P. Wonka, P. Müller, P. Waddell, and B. Watson, "Modeling the appearance and behavior of urban spaces," in *Eurographics 2009, State of the Art Report, EG-STAR*. Eurographics Association, 2009.
- [44] P. Waddell, "Urbansim: Modeling urban development for land use, transportation and environmental planning," *Journal of the American Planning Association*, vol. 68, pp. 297–314, 2002.
- [45] B. Watson, P. Müller, O. Vervovka, A. Fuller, P. Wonka, and C. Sexton, "Procedural urban modeling in practice," *IEEE Comput. Graph. Appl.*, vol. 28, no. 3, pp. 18–26, 2008.
- [46] B. Weber, P. Mueller, P. Wonka, and M. Gross, "Interactive geometric simulation of 4d cities," *Computer Graphics Forum*, April 2009.
- [47] L.-Y. Wei, S. Lefebvre, V. Kwatra, and G. Turk, "State of the art in example-based texture synthesis," in *Eurographics 2009, State of the Art Report, EG-STAR*. Eurographics Association, 2009.
- [48] L.-Y. Wei and M. Levoy, "Fast texture synthesis using tree-structured vector quantization," in *Proc. Of ACM SIGGRAPH '00*, 2000, pp. 479–488.
- [49] P. Wonka, M. Wimmer, F. Sillion, and W. Ribarsky, "Instant architecture," in *Proc. Of ACM SIGGRAPH '03*, 2003, pp. 669–677.
- [50] H. Zhou and J. Sun, "Terrain synthesis from digital elevation models," *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 4, pp. 834–848, 2007.
- [51] K. Zhou, X. Huang, X. Wang, Y. Tong, M. Desbrun, B. Guo, and H.-Y. Shum, "Mesh quilting for geometric texture synthesis," in *SIGGRAPH*, pp. 690–697, 2006.




Fig. 7. Because model synthesis is inefficient on curved models bounding volumes are used to simplify the geometry (a). The bounding boxes are cut into two objects, so the dome will scale uniformly and the cylinder will scale uniformly in x and y . The output is generated and the complex original shapes are substituted back in (b,c).




Fig. 8. Many complex buildings (b) are generated from four simple ones (a). The output contains many vertices that have been constrained to intersect four faces and a few of these vertices are circled. The result also uses the connectivity constraint to space the buildings apart which gives the buildings more room to develop into more interesting shapes.




Fig. 9. A fleet of spaceships (b,c) is automatically generated from a simple spaceship model (a). Without the connectivity constraint several dozen small unconnected spaceships are generated (b), but they are all packed together. With the connectivity constraint, six large spaceships are generated (c). Dimensional constraints are extensively used to ensure the rocket engines and other structures do not stretch unnaturally. The shape of the spaceships have a high genus because there are gaps in between the beams and parts of the spaceships.




Fig. 10. A large fully connected road network is generated (b) from a few streets using the connectivity constraint. The dimensions of the roads are also constrained.




Fig. 11. Several long roller coasters (b) are generated from one simple ones (a). Dimensional constraints are used to keep the track a certain width.




Fig. 12. Large-scale constraints are used to build spaceships in the shape of the letter ‘G’, rectangular buildings in the shape of the letter ‘P’, and buildings from Figure 8 in the shape of the letter ‘M’.