

Spec2Fab: A Reducer-Tuner Model for Translating Specifications to 3D Prints

by

Desai Chen

B.S. in Computer Science, Carnegie Mellon University (2011)

Submitted to the Department of Electrical Engineering and Computer Science


in partial fulfillment of the requirements for the degree of

Master of Science in Computer Science and Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2013



© Massachusetts Institute of Technology 2013. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
August 30, 2013

Certified by
Wojciech Matusik
Associate Professor of Electrical Engineering and Computer Science
Thesis Supervisor

Accepted by
Leslie Kolodziejski
Chairman, Department Committee on Graduate Students

Spec2Fab: A Reducer-Tuner Model for Translating Specifications to 3D Prints

by

Desai Chen

Submitted to the Department of Electrical Engineering and Computer Science
on August 30, 2013, in partial fulfillment of the
requirements for the degree of
Master of Science in Computer Science and Engineering

Abstract

Multi-material 3D printing allows objects to be composed of complex, heterogeneous arrangements of materials. It is often more natural to define a functional goal than to define the material composition of an object. Translating these functional requirements to fabricable 3D prints is still an open research problem. Recently, several specific instances of this problem have been explored (e.g., appearance or elastic deformation), but they exist as isolated, monolithic algorithms. In this research, I propose an abstraction mechanism that simplifies the design, development, implementation, and reuse of these algorithms. The solution relies on two new data structures: a *reducer tree* that efficiently parameterizes the space of material assignments and a *tuner network* that describes the optimization process used to compute material arrangement. I provide an application programming interface for specifying the desired object and for defining parameters for the *reducer tree* and *tuner network*. I illustrate the utility of my new framework by implementing several fabrication algorithms as well as demonstrating the manufactured results.

Thesis Supervisor: Wojciech Matusik

Title: Associate Professor of Electrical Engineering and Computer Science

Acknowledgments

This research is made possible with the help of many knowledgeable people. First, I would like to thank my advisor, Wojciech Matusik for many insightful suggestions and technical discussions during the course of the project. My colleagues David I.W. Levin, Piotr Didyk, and Pitchaya Sitthi-Amorn have come up with numerous great ideas and have helped with developing these ideas into practical solutions to the problems that I encounter. Thanks to them for their creative minds and many hours of labor. Thanks to my lab mates Kiril Vidimče and Jonathan Ragan-Kelley for reading early drafts of this thesis and providing valuable comments. I would like to thank Hanspeter Pfister, Sylvain Paris, Fredo Durand and Ilya Baran for their helpful suggestions as well as Bernd Bickel, Moritz Bächer and Miloš Hašan for providing software and data.

Thank you to my father Mingda and my mother Yin for their constant support and for inspiring me to conduct rigorous research. Finally, I must thank Jiaqi, my best companion for making my life a wonderful journey.

Contents

Acknowledgments	3
List of Figures	6
1 Introduction	8
2 Related Work	11
3 Design Goals	15
4 Data Structures	17
4.1 The Reducer Tree	17
4.2 Tuners and Tuning the Reducer Tree	20
4.3 Tuner Network	21
5 Process Configuration	23
5.1 Defining the Reducer Tree	23
5.2 Defining a <i>Tuner</i>	24
5.3 Binding <i>Tuners</i>	25
5.4 Establishing the <i>Tuner Network</i>	25
5.5 Process Use	26
6 Experiments	27
6.1 Spatially-varying Albedo	27
6.2 Heterogeneous Subsurface Scattering	29

6.3	Goal-based Caustics	32
6.4	Elastic Behavior	35
6.5	Combining Deformable Object and Spatially-varying Albedo	36
6.6	Combining Caustics and Spatially-varying Albedo	38
7	Discussion	39
8	Conclusion	41
	Bibliography	42

List of Figures

4-1	2D representations of the <i>geometry nodes</i> used in <i>reducer trees</i>	18
4-2	Two examples of <i>reducer trees</i> . Left: A tree for caustics effect. Right: a tree for subsurface scattering properties.	19
4-3	A diagram of the <i>tuner</i> workflow. Arrows indicate flow and type of information passed between the individual <i>tuner</i> components.	20
4-4	Two <i>toners</i> attached to a <i>reducer tree</i> . Each <i>tuner</i> is responsible for tuning the nodes in its attached subtree (denoted by shapes with similarly colored outlines).	21
4-5	One example of <i>tuner network</i> that is used in experiments.	21
5-1	Abstract interface for Node and its two subclasses.	24
6-1	The reducer-tuner model enables creating objects of arbitrary shapes with embedded textures.	28
6-2	<i>Reducer tree</i> for texture.	29
6-3	<i>Reducer tree</i> for heterogeneous subsurface scattering effect.	30
6-4	A marble chessboard with prescribed subsurface scattering properties produced by Hašan et al. [11]. The insets show the samples under thin line illumination, and the graph shows the convergence of tuners for 10 out of 100 scattering profiles used in the example. The error is measured by square distance between two profiles, each containing 400 coefficients.	30
6-5	A marble dome produced by the subsurface scattering <i>reducer tree</i>	31

6-6	3D printed lens arrays that produce a caustic image of Einstein. The algorithm is proposed by Papas et al. [22]. Below we show convergence plots for the microfacet optimizations. The portrait is available from the United States Library of Congress’s Prints and Photographs division, now in the public domain.	33
6-7	<i>Reducer tree</i> for facet caustics.	33
6-8	A smooth surface that produces a caustic image of Einstein. The algorithm is proposed by Finckh et al. [10]. Below we show convergence plots for the smooth surface optimization.	34
6-9	<i>Reducer tree</i> for smooth caustics.	34
6-10	<i>Reducer tree</i> for elastic behavior.	35
6-11	A 3D printed book with prescribed deformation behavior under load. The plot shows the error (in meters) as a function of iteration number of our branch and bound based <i>tuner</i> . The blue line shows the smallest error seen so far while the red line the <i>tuner</i> ’s progress exploring material subtrees [4].	36
6-12	A miniature of Earth with prescribed deformation behavior.	37
6-13	<i>Reducer tree</i> for combining texture and elastic behavior.	37
6-14	A textured smooth surface producing a designed caustics image under proper illumination.	38
6-15	<i>Reducer tree</i> for combining texture and smooth caustics.	38

Chapter 1

Introduction

3D printing receives a lot of attention as it aims to democratize fabrication. The ever expanding range of printing materials allows for fabrication of complex objects with spatially varying appearance, optical characteristics, and mechanical properties. One of the most important unsolved problems in this area is how to compute an object's material composition from a functional or behavioral description. I will call this process *specification to fabrication* translation (Spec2Fab). The goal of this work is to provide a convenient abstraction for specifying such translators. This is necessary to move past the current direct specification model of 3D printing.

Today, 3D printing of an object requires a material be directly specified for each voxel inside the object volume. This approach is fraught with difficulties. First, 3D printable models become specific to a single printer type, i.e., the models are built from materials provided by a given printer. Consider the inconvenience that would result from word processing documents being compatible with specific 2D printers. Second, working directly with printing materials rather than material properties is extremely challenging for users. Imagine the difficulty in finding the right combination of printing materials that would provide a specific color, stiffness, or refractive index.

My work is motivated by the recent research efforts in the computer graphics community to create specific instances of the translation process, for example, subsurface scattering [9, 11] or deformation properties [5]. However, each of these instances is a custom, monolithic solution which is difficult to extend, combine, or modify. The

main insight is that all these process instances share a similar structure. First, they rely on the ability to accurately simulate the physical properties of an object given its geometry and material assignment. They use this simulation within an optimization framework to search the space of all possible material assignments in order to find the one that best reproduces the desired properties. Due to the combinatorial nature of the search space the naive optimization approach is not tractable. For example, when the printing volume has N voxels and each of these voxels can be assigned to one of M base materials, the search space has N^M dimensions. To overcome this problem, the search space is reduced to a lower-dimensional space using a reduction model. The goal of the reduction step is to aggressively shrink the search space in a domain-specific manner such that it still contains good approximations to the optimal solution. This search space reduction combined with the right choice of the optimization algorithm delivers a computationally tractable approximation.

The reduction-optimization structure suggests that it is possible to provide a more general abstraction mechanism for translating 3D models to printer and material-specific representations. This key observation leads to my thesis statement:

Spec2Fab processes share a similar structure and often use a small set of common components.

I take the first step in building a unified model for Spec2Fab translation and provide a small set of exemplar components as building blocks. Towards this end, I propose two novel data structures designed to aid the fabrication process:

- The *reducer tree* is a tree-based data structure that allows us to parameterize the space of material assignments.
- The *tuner network* is a data structure for specifying the optimization process.

My solution also provides an API for specifying the desired object, setting up the simulation, and defining parameters for the *reducer tree* and *tuner network*. In general, my framework simplifies the construction of new computational fabrication algorithms. More specifically, different components of the process can be easily replaced and other components easily reused. Various optimization strategies can also be

explored with lower implementation burden. In order to show these advantages, I illustrate how existing computational design processes fit into this framework and how they can be combined. I demonstrate the results of these algorithms on a variety of different examples fabricated using 3D printers.

Chapter 2

Related Work

The new data structures draw ideas from previous work in rendering and optimization. The *reducer tree* is inspired by Cook’s shade trees [7] and their modern implementation in current rendering systems (e.g., Maya, RenderMan, etc.). Using these approaches, complex effects can be achieved by combining a set of basic shading blocks. The *reducer tree* also uses a tree data structure that combines a set of primitives to compute a material assignment for each point inside of an object volume and describes a spatial-partition of the object volume. While there are many space-partitioning data structures (BSPs, Quad-trees, KD-trees, etc.), they are difficult to specify by hand and they are typically tied to the object geometry. My *reducer tree* is intuitive to construct but sufficiently general for representing material distributions. In addition, it is specified in an object-independent manner – the same *reducer tree* can be reused for processing different objects. My second, new data structure which is responsible for optimizing material assignment is the *tuner network*. It is inspired by probabilistic graphical models [13] that have been popular for representing variable dependence in multi-variate optimization problems. These problems can be parallelized [18] according to the associated graph structure.

This work is also inspired by the many instances of specification to fabrication translation pipelines that have been proposed recently. These pipelines span a wide range of functional goals (e.g., mechanical, optical, appearance). They drive the simulation, optimization, and reductions chosen for each method. One example of such

processes is recent work on optimizing material composition to achieve prescribed deformation behavior. Bickel and colleagues [5] have designed a system for manufacturing multi-layer composites with a given elastic behavior. Skouras *et al.* [24] provide design and construction of balloons with a prescribed shape while inflated. Furthermore, material optimization has been employed to create mechanical clones of human faces [6].

Optimizing and manufacturing objects with desired appearance and optical characteristics has also been explored. Weyrich *et al.* [27] compute surface micro-geometry that yields a desired BRDF. Similar approaches have been proposed [10, 22] to produce refractive surfaces that form user-defined caustics. These methods have been extended to optically decrypt hidden images [23]. Complementary work examines fabricating surfaces with spatially varying reflectance [19, 20] and diffuse shading [1]. Another set of approaches uses optimization to compute shadow casting surfaces and volumes [2, 3, 21] that reproduce a given set of input images. Finally, optimization-based approaches have also been employed to control the subsurface scattering of 3D printed multi-layered models [9, 11]. I seek to exploit the common form of the above works in order to generalize them.

There have been studies of material assignment representations in other fields, primarily in mechanical engineering. Here I only list a few representative works. Kumar *et al.* [17] describe material composition by dividing a volume into sub-volumes. They perform material interpolation using local, sub-volume coordinate systems. Jackson [12] explores several mesh data structures for spatial sub-division. Kou and Tan [15] give a comprehensive review of spatial partition schemes and material interpolation functions. Kou *et al.* [16] use a hierarchy of procedures to define material composition with a small number of design parameters. They run particle swarm optimization [14] on the design parameters to minimize thermal stress of an object. Their work focused on smoothly varying materials. There is no discussion of high-frequency discrete material assignment for capturing details. They are limited to global optimization algorithms such as Particle Swarm Optimization because the dependency structure between the design variables is not modeled. In computer graphics, procedural de-

scriptions for material assignment have also been studied. Cutler *et al.* [8] describe a scripting language for specifying layered solid models and describing material composition. Vidimče *et al.* [26] provide a fabrication language and a programmable pipeline for specifying material composition directly and precisely throughout a volume. In contrast, Spec2Fab can be used to design algorithms that only require object properties rather than a precise description of material assignment.

In this research I focus on constructing a common model for encompassing the entire Spec2Fab problem. I am motivated by the similarities present in state-of-the-art algorithms. Table 2.1 shows the coordinate reduction methods as well as the optimization schemes used in all these prior approaches. A key observation is that previous works share a common methodology and a common set reduced coordinate components. These components are then combined with (often off-the-shelf) optimization techniques to form new material assignment algorithms. This suggests that one could find a small set of components that can be tied together to synthesize advanced methods. I attempt to identify these components as well as a suitable model for their interaction.

Paper	Goal	Reduced Coordinates	Optimization
[10]	Optical (Caustics)	B-Spline Surface	SPSA ¹
[22]	Optical (Caustics)	Piecewise Constant Tiles	SA ²
[1]	Optical (Relief)	Height Field	Gradient Descent
[27]	Optical (Reflectance)	Piecewise Constant Tiles	SA
[23]	Optical (Refraction)	Piecewise Constant Tiles	SA
[21]	Optical (Shadows)	Voxel Grid	Custom Discrete
[2]	Optical (Shadows)	Layered Materials	Quadratic Program
[3]	Optical (Shadows)	Height Field	Custom/SA
[9]	Optical (Subsurface)	Layered Materials	Conjugate Gradient
[11]	Optical (Subsurface)	Layered Materials	Branch and Bound
[5]	Mechanical (force)	Layered Materials	Branch and Bound
[6]	Mechanical (shape)	Height Field	Newton-Raphson
[24]	Mechanical (shape)	Triangle Mesh	ALM ³

Table 2.1: The goal type, reduction type and optimization used by prior computational fabrication approaches.

¹Simultaneous perturbation stochastic approximation (SPSA).

²Simulated annealing (SA).

³Augmented Lagrangian method (ALM).

Chapter 3

Design Goals

The design of my general translation framework is guided by the following principles:

- **Modularity:** Spec2Fab translators are complicated both algorithmically and from a software engineering point of view. To combat this, any proposed framework must break the problem into a manageable number of small, reusable building blocks.
- **Extensibility:** Developers must be able to add their own building blocks to the system. This allows the system to grow in conjunction with the capabilities of newer 3D printers.
- **Device Independence:** Spec2Fab translators should be device independent. They should be easily adaptable to different types of 3D printers.
- **Input Geometry Independence:** Spec2Fab translators should be geometry independent. For example, a process for applying a texture to a 3D printed object should work for **any** object.

I separate the Spec2Fab process into two phases, the process configuration phase and the process use phase. The process configuration phase is typically done once by a skilled developer who constructs a Spec2Fab translator. The process use phase is typically performed multiple times by an end-user who is only required to provide an

object specification (e.g., object geometry and deformation properties) and a target device.

The process configuration phase produces a Spec2Fab translator which will assign a desired volumetric material distribution to a user supplied input geometry given a user specified goal. A developer can describe this phase using two new data structures, the *reducer tree* and *tuner network*. The *reducer tree* parameterizes a volumetric material assignment using a small set of *geometry* and *material nodes* while the *tuner network* is used to describe an optimization process as a connected set of *tuner* objects. Both *nodes* and *tuners* can be easily recombined and reused thus making my framework highly modular. Furthermore, *nodes* and *tuners* define abstract interfaces and thus developers can easily add new types of each. This makes my framework extensible.

Reducers and *tuners* are chosen to be independent of printer capabilities. Instead, a user can account for printer type by altering the available materials that the translator may assign to the input geometry. This is important since it grants the *reducer-tuner* model device independence. Finally, the *geometry nodes* of the *reducer tree* are designed to function irrespective of input shape. Since the Spec2Fab translators represented by my data structures use a composition of these node types, they are geometry independent.

In the following chapters I will describe the *reducer tree* and the *tuner network*, their constituent components and the mechanisms by which they interact.

Chapter 4

Data Structures

In this chapter, I describe the structure of the *reducer tree* as well as all types of *reducer nodes* that are implemented in my framework. I also show the structure of individual *tuners* and how they can be arranged into the *tuner network*.

4.1 The Reducer Tree

Estimating material assignments at output device resolution is computationally intractable. Therefore, material assignment has to be computed using a reduced representation. A developer specifies this representation using a *reducer tree*. This structure is conceptually similar to those used in programmable shading systems (such as Cook's shade trees [7] or Maya Shader Networks). These systems are primarily concerned with assigning known materials and textures to an object's surface whereas I am seeking an optimal volumetric assignment from a defined set of materials. In order to accomplish this task, a developer builds a tree-based data structure which contains the entire object volume as its root node. I define two classes of *reducer nodes*: *geometry nodes* and *material nodes*.

Geometry Nodes: A *geometry node* takes a volumetric region as input and produces a partition of this region into smaller sub-regions. To demonstrate the flexibility of the *reducer tree*, I define a small yet powerful set of partitioning nodes which are

described as follows (Figure 4-1):

- Plane Node – partitions the space into two half-spaces,
- Column Node – takes as input the number of columns and accordingly partitions the space,
- Voxel Node – takes as input voxel size and uniformly partitions the space,
- B-spline Node – partitions a volume into two regions cut by a B-spline,
- Stratum Node – takes a single positive distance parameter as an input and partitions the volume into two regions divided by the iso-distance surface.




Figure 4-1: 2D representations of the *geometry nodes* used in *reducer trees*.

Material Nodes: The *leaf nodes* of the *reducer tree* are the material parameterization nodes. These contain a material assignment function $\lambda(x)$ that maps spatial coordinates to materials. *Material nodes* can assign a void material to regions of the volume. Hence, surface displacements and material assignments can be treated in a unified fashion. While *material node* can be extended to implement arbitrary material assignment functions (such as functionally graded materials and lattice structures), all the following results require only a single type of *material node*, *layered node*, which assigns k material layers of varying thickness to a geometry partition. As a special case, a constant material for a region can be implemented as a *layered node* with a single layer.

Geometry nodes and *material nodes* can be connected into a tree structure that describes a material assignment function throughout the input geometry. The essential property of this data structure is that it naturally adapts to the input geometry. This key feature allows it to be reused for different shapes. Figure 4-2 shows two

examples of *reducer trees* and the resulting geometric distributions of materials. The tree on the left can be used to produce objects with caustics effects. The input shape is first divided into columns. Next, each column is sliced by a plane. By assigning materials to lower parts of the columns, the microfacet surface is created. On the right, I show a tree for producing objects with subsurface scattering properties. The input shape is first divided using a *stratum* node. Then, the outer layer is sliced into columns, which create a texture once materials are assigned. The inner part of the object obtains a material with given subsurface scattering properties.

Since the tree describes a nested set of partitions, it can efficiently perform material queries by passing a query point through the tree until it arrives at a parameterization node. Note that the *reducer tree* does not inherently enforce material continuity between disjoint regions of the object; however, for material assignment problems, this is neither required nor desired.

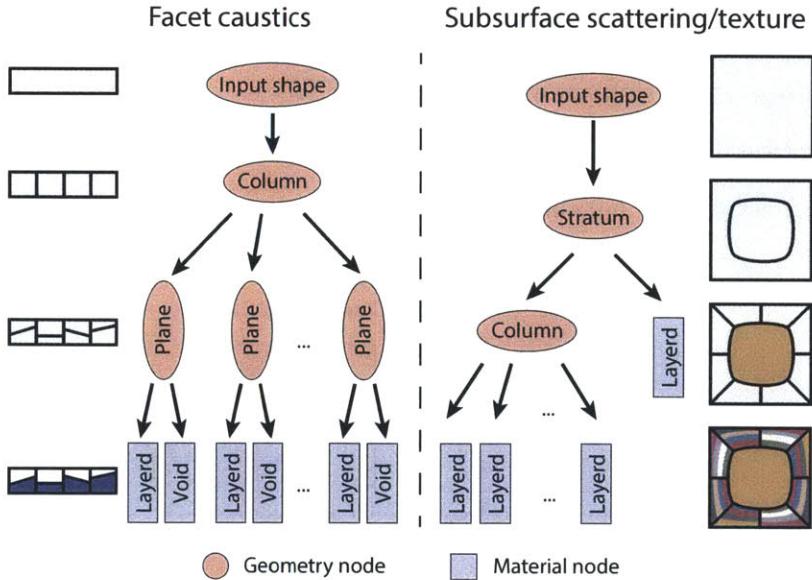


Figure 4-2: Two examples of *reducer trees*. Left: A tree for caustics effect. Right: a tree for subsurface scattering properties.

4.2 Tuners and Tuning the Reducer Tree

I define *tuners* as an abstract interface between a *reducer node* and an optimization algorithm. A *reducer node* has parameters that control its space partitioning and material composition, e.g., *stratum node* has a parameter which determines its thickness, *column node* has parameters that control the width of columns. A *tuner* is responsible for tuning these parameters to achieve a specified goal. A *tuner* is comprised of an optimization scheme, an error metric, a simulator and a goal (Figure 4-3). In order to create an optimal material assignment, *tuner nodes* must be attached to the *reducer tree*. The developer links each *tuner* to a *reducer node* in the *reducer tree*. Each *tuner* then traverses the *reducer subtree* (rooted at this *reducer node*) and constructs a parameter list by querying each visited *reducer node*. The developer can label parameters as fixed or free. *Tuner nodes* contain specific optimization routines (e.g., a quadratic program solver). The *tuner node* then uses this routine to optimize its associated free parameters. During execution, additional information (e.g., parameters, errors, etc.) from neighboring Tuner Nodes can be obtained via the *tuner network* (Section 4.3). Tuner execution can be scheduled (again by the programmer) allowing both serial and parallel processing.




Figure 4-3: A diagram of the *tuner* workflow. Arrows indicate flow and type of information passed between the individual *tuner* components.

Figure 4-4 shows a simple example in which two *tuners* are attached to sibling nodes in a *reducer tree*. The first *tuner* is responsible for *layered* and *void material nodes* while the second is responsible for several *layered material nodes*.

Tuner-reducer node connections




Figure 4-4: Two *tuners* attached to a *reducer tree*. Each *tuner* is responsible for tuning the nodes in its attached subtree (denoted by shapes with similarly colored outlines).

4.3 Tuner Network

For many fabrication problems, *tuners* should not act in isolation. For instance, an optimal material assignment at a given point depends on the material assignment in the neighborhood. Tuning the free parameters without taking into account this dependency usually yields suboptimal results. Therefore, I allow the *tuners* to share information according to a user-specified graph structure. This allows my framework to interface with a wide range of existing graph-based optimization and inference algorithms. For example, Figure 4-5 shows one of the *tuner networks*.

Tuner-tuner node connections

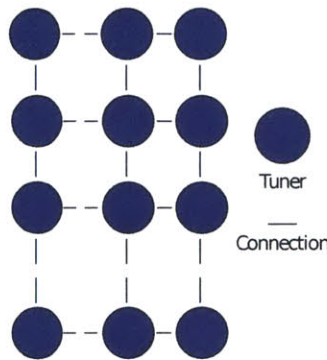


Figure 4-5: One example of *tuner network* that is used in experiments.

In this network, each node is connected to its four neighbors. In the following chapters, *tuners* enclosed in dashed boxes are connected in this form. *Tuner networks* do not always exhibit this regular connectivity. For instance, *tuners* can be completely unconnected (Figure 4-4) or be organized into groups which feature intra- but not inter-group connections.

Chapter 5

Process Configuration

In this section I show each step of process configuration. At the same time I describe my implementation of the *reducer tree* and *tuner network* data structures.

5.1 Defining the Reducer Tree

A *reducer tree* can be constructed expediently using existing *geometry* and *material node* types. *Reducer trees* are not restricted to use only existing node types. New node types can be added by extending the `ReducerNode` class (see Figure 5-1). In particular, implementing a new *geometry node* type requires providing the function `getOutputIndex`. This function takes, as input, a 3D position and returns the `id` of its child (`Geometry` or `Material`) that contains this 3D point. This function also computes a local coordinate for this child node. Performing computations in local coordinates allows us to abstract away the geometry of a given object. Implementing a new `MaterialNode` type requires specifying the `getMaterial` function which takes a 3D point and returns the material at this point in the local geometry coordinate system. Both types of nodes have an `evaluate` function which is responsible for updating their internal states. This function is used by the *tuner network* to modify the internal state of the nodes. As an example, we show a *reducer tree* for performing texture mapping (see Figure 4-2). We also provide the corresponding pseudo-code (see Algorithm 1).

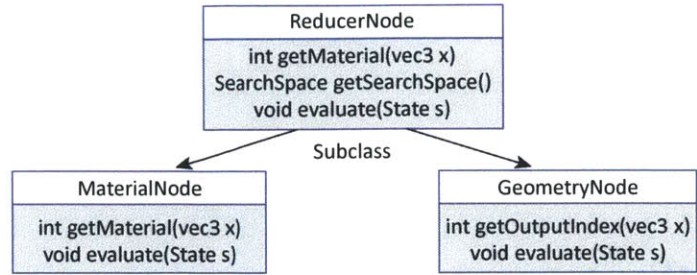


Figure 5-1: Abstract interface for Node and its two subclasses.

1. Create a root node **R** from `inputMesh`
2. Subdivide **R** into outer layer **O** and inner volume **V** (Stratum Node)
3. Subdivide **O** into set of columns **C** (Column Node)
4. For each column *c* in **C**
5. Subdivide *c* into two layers (Layer Node)
6. End

Algorithm 1: Constructing a *reducer tree* for texturing

5.2 Defining a *Tuner*

Recall that a *tuner* consists of four components: a simulation, an error metric, an optimizer and a goal (Figure 4-3). Certain combinations of goal, metric and simulator are not compatible (i.e., a deformation simulator is not compatible with an error metric that compares images). Our API checks and prevents such incompatible combinations of components. The optimization algorithm can request the error value for a given state using a callback function `getError()` which is defined by the *tuner*. Additional callback functions can be defined by the developer depending on the needs of the algorithm. For example, the branch and bound algorithm requires a custom function to compute error bounds for a given state.

5.3 Binding *Tuners*

Tuners are assigned to nodes in the *reducer tree* using the `setNode` function. Once assigned, a *tuner* can optimize the parameters of its associated subtree. *Reducer nodes* provide a `getSearchSpace` function which returns all free variables in the node subtree. In order to make *tuners* as flexible as possible we provide a `Parameter` class. Parameters can be either discrete or continuous, they can have associated bounds, and they can be marked as free or fixed.

5.4 Establishing the *Tuner Network*

The *tuner network* is an undirected graph that describes connections between *tuners*. *Tuner nodes* store a list of their neighbors. Only neighboring *tuners* are allowed to exchange information. In our current implementation, this is accomplished using a shared memory array. As an example, we show how to construct and initialize a *tuner network* for a simple optimization scheme – a Simulated Annealing algorithm [25] (see Algorithm 2). The *tuner network* also requires a schedule that specifies in what order individual *tuners* should be executed. This schedule is specified by the developer. Once the *tuner network* is constructed, the process configuration phase is complete. We obtain a compiled executable that computes desired material assignment from an input specification.

```

1.  for each Tuner  $T_i$  attached to a plane
2.    for each Tuner  $T_j$  adjacent to  $T_i$ 
3.      add  $T_j$  to  $T_i$ 's list of neighbors
4.    end
5.  end
6.  iterate N times
7.  for each Tuner  $T_i$ 
8.    set temperature for  $T_i$ 's optimizer
9.    run  $T_i$ 
10. end

```

Algorithm 2: Connecting and executing the *tuner network*

5.5 Process Use

The compiled program, which executes the *tuner network*, takes five types of arguments: the input geometry, the goal, the simulation configuration, a set of materials and the target 3D printer specification. After the *tuner network* is executed, the parameters of the *reducer nodes* in the *reducer tree* are set. It is then straightforward to compute the material assignment at arbitrary resolution. Since, a typical multi-material 3D printer requires a volumetric model with per voxel material assignment, we can simply iterate over all voxels in the volume. We obtain the material assignment by evaluating (`getMaterial`) of the *reducer tree* root at the center of each voxel location. This representation can be easily converted to a printer specific format for output. For the Objet500 Connex printer used in this paper we extract material isosurfaces which are submitted to the printer as STL files.

Chapter 6

Experiments

In order to evaluate capabilities of our system, I have implemented a number of existing translation processes. Furthermore, the ease with which different algorithm components can be combined enables the creation of two new translation processes. The first combination is an algorithm that can produce objects with desired refractive properties and an associated texture. The second algorithm applies a desired texture to an object with prescribed deformation behavior. All of these processes, both from prior work and new ones, should easily fit into our framework. I printed the objects using a Stratasys Object500 Connex multi-material printer. The following paragraphs provide a detailed description of how the individual algorithms can be designed with my system.

6.1 Spatially-varying Albedo

I have designed a Spec2Fab translator that allows 3D printing of textured models (Figure 6-1). The *reducer trees* and the *tuner networks* for producing these results are presented in Figure 6-2. Being able to apply precisely specified spatially-varying albedo values to printed models is a crucial capability. However, no standard processes have been designed for this task so far. The input to the texturing algorithm is a shape and its desired albedo texture. Since the texture is only affected by materials close to the surface, I use a *stratum node* to divide the input shape into a thin shell




Figure 6-1: The reducer-tuner model enables creating objects of arbitrary shapes with embedded textures.

and an inner volume. I then divide the outer layer into columns. The set of printable colors is expanded by stacking translucent materials using the `LayeredMaterial` Node. The number of layers is fixed to the number of print materials. The reduced parameters are the thickness values of each material layer. For each column, the *tuner's* optimizer looks up the proper stacking that produces the closest albedo value. Due to printer resolution, the range of albedo values that can be achieved is quantized. I therefore implement an error diffusion algorithm by connecting neighboring *tuners*. In this simple algorithm, the simulation is a table lookup using measured albedo value corresponding to different base materials.




Figure 6-2: *Reducer tree* for texture.

6.2 Heterogeneous Subsurface Scattering

I have replicated the subsurface scattering process of Hašan et al. [11] using a tree shown in Figure 6-3. A 3D printed chessboard is shown in Figure 6-4. The input to this algorithm is a 3D mesh along with subsurface-scattering profiles defined at a set of surface points. We use the same *reducer tree* as in the texture example thereby simplifying the process configuration phase. Since the *reducer tree* adapts to the input geometry, I can apply the same marble material to arbitrary meshes such as the dome shown in Figure 6-5. The only difference is that I allow each column to have four layers of varying thickness and material. I use a branch and bound optimization algorithm which has been modified to handle continuous parameters by allowing discrete increments. I implement a bound estimate callback function specific to this problem. Each column is optimized independently using the algorithm. The simulation computes a scattering profile for a given stacking and the error metric compares the simulated and goal profiles using squared error.

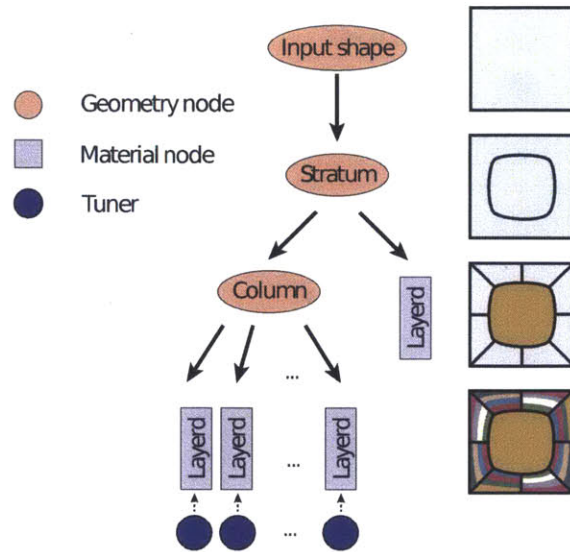


Figure 6-3: *Reducer tree* for heterogeneous subsurface scattering effect.




Figure 6-4: A marble chessboard with prescribed subsurface scattering properties produced by Hašan et al. [11]. The insets show the samples under thin line illumination, and the graph shows the convergence of tuners for 10 out of 100 scattering profiles used in the example. The error is measured by square distance between two profiles, each containing 400 coefficients.




Figure 6-5: A marble dome produced by the subsurface scattering *reducer tree*.

6.3 Goal-based Caustics

I have configured two different processes for computing goal-based caustics. While they define exactly the same goal, they have very different *reducer trees* and *tuner networks*. The first process is based on the work of Papas et al. [22]. It computes a set of micro-lenses which produce the desired caustic image as shown in Figure 6-6. The image is pre-processed into a set of Gaussian distributions. Each distribution is matched with a micro-lens. The optimization applies simulated annealing to permute the location of these micro-lenses in order to construct a smooth surface. The micro-lenses are represented using *plane nodes*. The complete reducer tree is shown in Figure 6-7. In the *tuner network*, each *tuner* is connected to its four neighbors. During the execution of an individual *tuner*, the optimization algorithm makes a randomized decision about whether or not to swap its micro-lens with one of its neighbors based on smoothness of the surface. The *tuners* are executed many times until a user-specified convergence criterion is met.

The second process is based on the work of Finckh et al. [10] (Figure 6-8). I use a *B-spline node* to represent a smooth surface (Figure 6-9). This is in contrast to the potentially discontinuous surface in the method above. The reduced parameters are the height of each spline control point. I implement a simple caustics simulator for height fields. The simulated image is compared to the goal image using mean squared error.

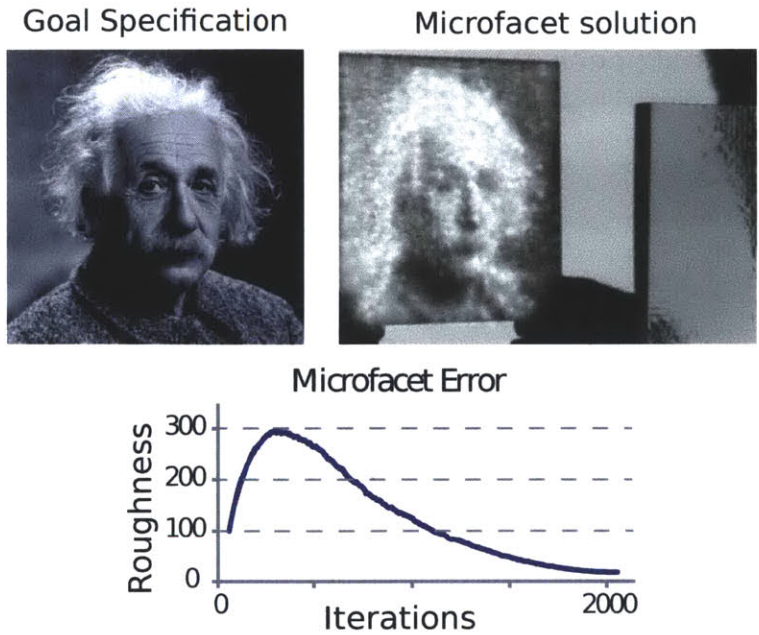


Figure 6-6: 3D printed lens arrays that produce a caustic image of Einstein. The algorithm is proposed by Papas et al. [22]. Below we show convergence plots for the microfacet optimizations. The portrait is available from the United States Library of Congress’s Prints and Photographs division, now in the public domain.




Figure 6-7: *Reducer tree* for facet caustics.

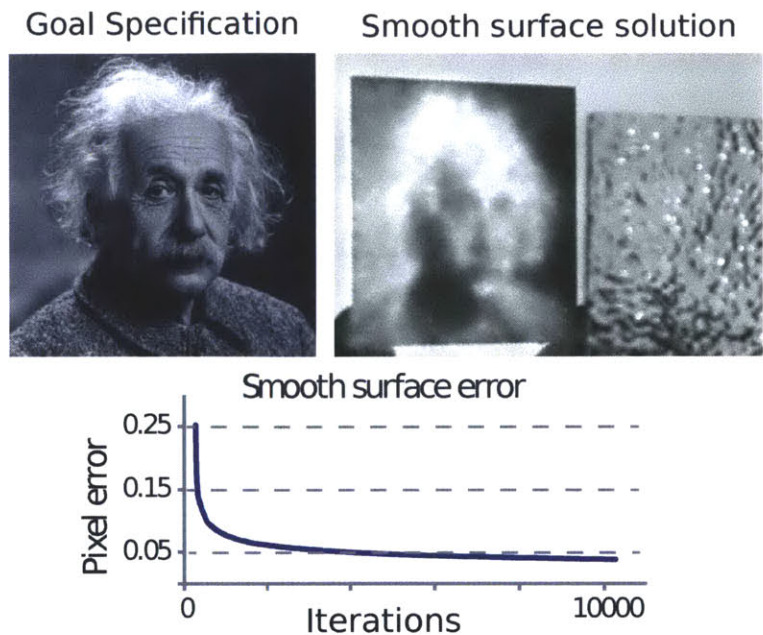


Figure 6-8: A smooth surface that produces a caustic image of Einstein. The algorithm is proposed by Finckh et al. [10]. Below we show convergence plots for the smooth surface optimization.

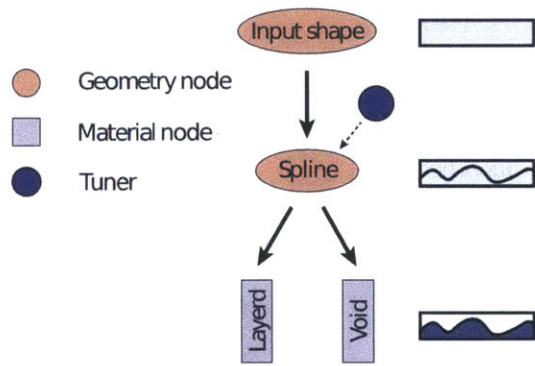


Figure 6-9: *Reducer tree* for smooth caustics.

6.4 Elastic Behavior

In the spirit of Bickel et al. [4], I have implemented an algorithm to compute material distribution based on a desired force-displacement response. The input to this algorithm is a mesh, a simulation configuration, and a desired shape. Simulation configuration includes vertex constraints and forces applied to the mesh. For this example, I use a co-rotational finite element method (FEM) simulation with linearly elastic materials to estimate the objects deformation. For the *reducer tree* (Figure 6-10), I use a voxel partition to divide the object into a low-resolution grid. I assign a single material to each grid cell. The FEM simulator queries the *reducer tree* for material assignments at arbitrary spatial locations. I use the same branch and bound algorithm as in our subsurface-scattering process but with a different bound computation callback function. I use the mean squared distance between the simulated and the desired shapes as the error metric. I have designed a simple experiment to validate this process in which I set the goal of our optimization to be a given deformed state (Figure 6-11).

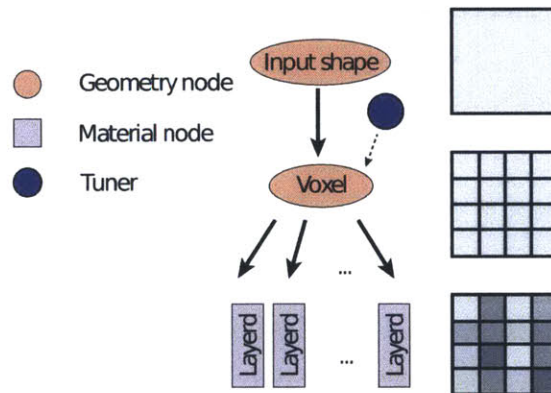


Figure 6-10: *Reducer tree* for elastic behavior.

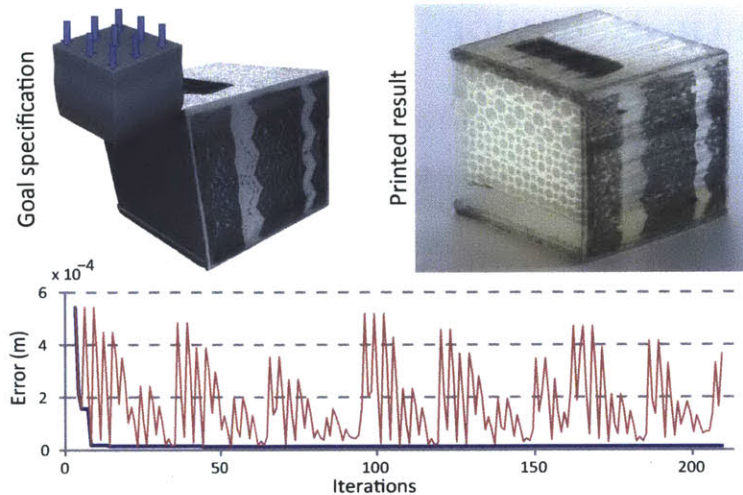


Figure 6-11: A 3D printed book with prescribed deformation behavior under load. The plot shows the error (in meters) as a function of iteration number of our branch and bound based *tuner*. The blue line shows the smallest error seen so far while the red line the *tuner*'s progress exploring material subtrees [4].

6.5 Combining Deformable Object and Spatially-varying Albedo

The first of my new translation processes combines spatially-varying albedo and elastic deformation properties (Figure 6-12). This is a very useful combination since when modeling objects we would like to specify both their appearance and “feel”. In this process, the input shape is divided into a thin shell and an inner volume (Figure 6-13). I optimize for the deformation behavior and the texture independently due to the limitations of our current FEM simulator. As the outer shell is very thin, it has negligible influence on overall object deformation.




Figure 6-12: A miniature of Earth with prescribed deformation behavior.




Figure 6-13: *Reducer tree* for combining texture and elastic behavior.

6.6 Combining Caustics and Spatially-varying Albedo

My second new combined translation process incorporates both smooth caustics and texture mapping (Figure 6-14). More specifically, I compute a transparent slab with a texture image that, when illuminated, casts a prescribed caustic image. The input slab is split into two pieces using a *plane node* as shown in Figure 6-15. The top piece is tuned to produce an input image. The material in the top piece is then fixed. The bottom piece is then tuned to produce a caustics image.




Figure 6-14: A textured smooth surface producing a designed caustics image under proper illumination.

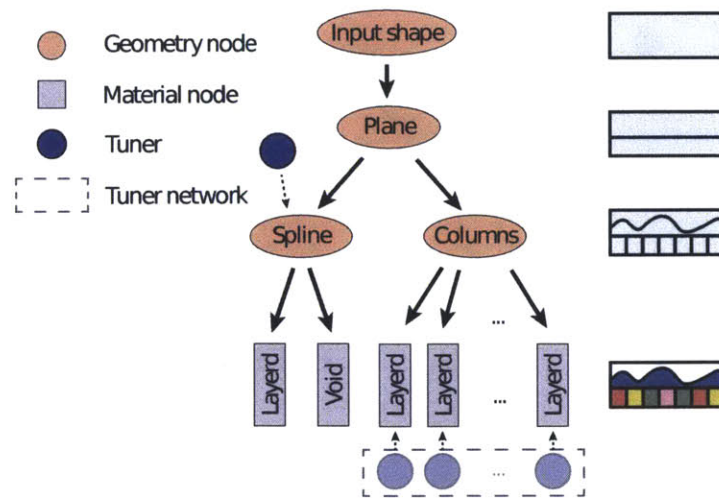


Figure 6-15: *Reducer tree* for combining texture and smooth caustics.

Chapter 7

Discussion

Through my series of experiments, The *reducer tree* and *tuner network* have allowed me to reuse a large number of software components. In the extreme case, I arrive at my subsurface scattering algorithm by trivially adapting the same *reducer tree* for texture mapping objects. The power of component reuse is further elucidated by the presence of *column nodes* in most examples and by the reuse of optimization schemes (such as branch and bound) in multiple algorithms. The *reducer tree* and *tuner network* make these similarities easy to observe and exploit. The only component whose reusability was not demonstrated in this proposal is simulation. In this work, I have aimed to fabricate objects with a wide range of physical properties and this necessitates the use of a wide range of simulation algorithms. Finally, once a whole process is configured, it is independent of input geometry and goal parameters. For example, I have run the spatially-varying albedo process with different geometries and input textures.

The *reducer-tuner model* is ideally suited for multi-material printing capable of producing objects with a wide range of different properties. In order to showcase these strengths, I have fabricated my examples for an Objet500 Connex – a phase change inkjet printer that uses photopolymers with a wide range of optical and mechanical properties. Even though only two materials can be used and mixed within a single object, my framework has been proven to be very useful. As the number of materials that can be printed simultaneously will grow, I expect the methodology presented

here to increase in utility. For other types of 3D printing technologies, Spec2Fab framework has a reduced use. In the case of 3D printing using a single, rigid material (e.g., fused filament fabrication or Stereolithography) the framework can be used to tune geometry of objects. For plaster-based 3D printers that produce full-color 3D prints, Spec2Fab framework can be employed to compute proper texturing of objects.

Chapter 8

Conclusion

In this research I have taken the first step towards solving an open problem in computational fabrication – creating a general translation process that transforms user-defined model specifications into printer and material-specific representations. My process relies on two data structures to make this general translation process expressive and computationally tractable: a *reducer tree* and a *tuner network*. I have shown how existing instances of this translation can be expressed and combined within my system. I believe that my API and its reference implementation will simplify and encourage development of new translation processes.

My system offers many exciting opportunities for future work. First, it would be extremely useful to implement many additional simulators in order to allow computing a variety of other properties, e.g., structural soundness, stability, material cost, and printing time. These simulators could be employed to expand the range of possible user-defined specifications. Similarly, only relatively simple error metrics have been proposed within the tuning process. The development of more sophisticated and, in particular, perceptually-driven material metrics remains a relatively unexplored research area. Finally, it would be very beneficial to couple my API with a visual interface to further simplify the task of translator construction. It is not obvious what the best visual interface for specifying functional or physical properties of these objects is and more research in this area is necessary.

Bibliography

- [1] Marc Alexa and Wojciech Matusik. Reliefs as images. *ACM Trans. on Graphics (SIGGRAPH 2010)*, 29(4):60:1–60:7, July 2010.
- [2] Ilya Baran, Philipp Keller, Derek Bradley, Stelian Coros, Wojciech Jarosz, Derek Nowrouzezahrai, and Markus Gross. Manufacturing layered attenuators for multiple prescribed shadow images. *Computer Graphics Forum*, 31(2):603–610, May 2012.
- [3] Amit Bermano, Ilya Baran, Marc Alexa, and Wojciech Matusik. ShadowPIX: Multiple images from self-shadowing. *Computer Graphics Forum*, 31(2):593–602, May 2012.
- [4] Bernd Bickel, Moritz Bächer, Miguel A. Otaduy, Wojciech Matusik, Hanspeter Pfister, and Markus Gross. Capture and modeling of non-linear heterogeneous soft tissue. *ACM Trans. on Graphics (SIGGRAPH 2009)*, 28(3):89:1–89:9, July 2009.
- [5] Bernd Bickel, Moritz Bächer, Miguel A. Otaduy, Hyunho Richard Lee, Hanspeter Pfister, Markus Gross, and Wojciech Matusik. Design and fabrication of materials with desired deformation behavior. *ACM Trans. on Graphics (SIGGRAPH 2010)*, 29(4):63:1–63:10, July 2010.
- [6] Bernd Bickel, Peter Kaufmann, Mélina Skouras, Bernhard Thomaszewski, Derek Bradley, Thabo Beeler, Phil Jackson, Steve Marschner, Wojciech Matusik, and Markus Gross. Physical face cloning. *ACM Trans. on Graphics (SIGGRAPH 2012)*, 31(4):118:1–118:10, July 2012.

- [7] Robert L. Cook. Shade trees. *Computer Graphics (SIGGRAPH 84)*, 18(3): 223–231, Jan 1984.
- [8] Barbara Cutler, Julie Dorsey, Leonard McMillan, Matthias Müller, and Robert Jagnow. A procedural approach to authoring solid models. *ACM Trans. on Graphics (SIGGRAPH 2002)*, 21(3):302–311, July 2002.
- [9] Yue Dong, Jiaping Wang, Fabio Pellacini, Xin Tong, and Baining Guo. Fabricating spatially-varying subsurface scattering. *ACM Trans. on Graphics (SIGGRAPH 2010)*, 29(4):62:1–62:10, July 2010.
- [10] Manuel Finckh, Holger Dammertz, and Hendrik P. A. Lensch. Geometry construction from caustic images. *Computer Vision (ECCV 2010)*, 6315:464–477, 2010.
- [11] Miloš Hašan, Martin Fuchs, Wojciech Matusik, Hanspeter Pfister, and Szymon Rusinkiewicz. Physical reproduction of materials with specified subsurface scattering. *ACM Trans. on Graphics (SIGGRAPH 2010)*, 29(4):61:1–61:10, July 2010.
- [12] Todd Robert Jackson. *Analysis of Functionally Graded Material Object Representation Methods*. PhD thesis, Massachusetts Institute of Technology, 2000.
- [13] Michael I. Jordan, Zoubin Ghahramani, Tommi S. Jaakkola, and Lawrence K. Saul. An introduction to variational methods for graphical models. *Machine Learning*, 37:183–233, 1999. ISSN 0885-6125.
- [14] James Kennedy and Russell Eberhart. Particle swarm optimization. In *Neural Networks, 1995. Proceedings., IEEE International Conference on*, volume 4, pages 1942–1948. IEEE, 1995.
- [15] X.Y. Kou and S.T. Tan. Heterogeneous object modeling: A review. *Computer-Aided Design*, 39(4):284 – 301, 2007.

- [16] XY Kou, GT Parks, and ST Tan. Optimal design of functionally graded materials using a procedural model and particle swarm optimization. *Computer-Aided Design*, 44(4):300–310, 2012.
- [17] V. Kumar, D. Burns, D. Dutta, and C. Hoffmann. A framework for object modeling. *Computer-Aided Design*, 31(9):541 – 556, 1999. ISSN 0010-4485.
- [18] Yucheng Low, Joseph Gonzalez, Aapo Kyrola, Danny Bickson, Carlos Guestrin, and Joseph M. Hellerstein. Graphlab: A new parallel framework for machine learning. *Conference on Uncertainty in Artificial Intelligence (UAI)*, July 2010.
- [19] Tom Malzbender, Ramin Samadani, Steven Scher, Adam Crume, Douglas Dunn, and James Davis. Printing reflectance functions. *ACM Trans. on Graphics (SIGGRAPH 2012)*, 31(3):20:1–20:11, May 2012.
- [20] Wojciech Matusik, Boris Ajdin, Jinwei Gu, Jason Lawrence, Hendrik P. A. Lensch, Fabio Pellacini, and Szymon Rusinkiewicz. Printing spatially-varying reflectance. *ACM Trans. on Graphics (SIGGRAPH Asia 2009)*, 28(5):128:1–128:9, Dec. 2009.
- [21] Niloy J. Mitra and Mark Pauly. Shadow art. *ACM Trans. on Graphics (SIGGRAPH Asia 2009)*, 28(5):156:1–156:7, Dec 2009.
- [22] Marios Papas, Wojciech Jarosz, Wenzel Jakob, Szymon Rusinkiewicz, Wojciech Matusik, and Tim Weyrich. Goal-based caustics. *Computer Graphics Forum*, 30(2):503–511, 2011.
- [23] Marios Papas, Thomas Houit, Derek Nowrouzezahrai, Markus Gross, and Wojciech Jarosz. The magic lens: refractive steganography. *ACM Trans. on Graphics (SIGGRAPH Asia 2012)*, 31(6):186:1–186:10, November 2012.
- [24] Mélina Skouras, Bernhard Thomaszewski, Bernd Bickel, and Markus Gross. Computational design of rubber balloons. *Computer Graphics Forum*, 31(2): 835–844, May 2012.

- [25] P.J. van Laarhoven and E.H. Aarts. *Simulated annealing: theory and applications*, volume 37. Springer, 1987.
- [26] Kiril Vidimče, Szu-Po Wang, Jonathan Ragan-Kelley, and Wojciech Matusik. OpenFab: A programmable pipeline for multi-material fabrication. *ACM Trans. on Graphics (SIGGRAPH 2013)*, 32(4), July 2013.
- [27] Tim Weyrich, Pieter Peers, Wojciech Matusik, and Szymon Rusinkiewicz. Fabricating microgeometry for custom surface reflectance. *ACM Trans. on Graphics (SIGGRAPH 2009)*, 28(3):32:1–32:6, July 2009.