Distributed Design and Fabrication of Parts with Local Composition Control

H. Wu, E. Sachs, N. Patrikalakis, D. Brancazio, J. Serdy, T. R. Jackson, W. Cho, H. Liu, M. Cima
Massachusetts Institute of Technology, Cambridge, MA

R. Resnick
ExtrudeHone Corporation, Irwin, PA

Abstract

Solid Freeform Fabrication (SFF) technologies offer the potential to create parts with local control of the composition (LCC). This work develops the representations, methods, tools and design rules needed to allow for the design and fabrication of parts with locally varying composition. Three Dimensional Printing is used as a prototypical Solid Freeform Fabrication process. The goal is to provide the designer with tools sufficient to allow the designer to create a successful part with "one-way information flow", that is, flow of information from the designer through the fabrication sequence without iteration. This paper focuses on the post-processing of representations, which includes information on local composition control. Such representations are first rendered as halftone images, in accordance with the digital aspect of Three Dimensional Printing. The boundary geometry information is then reconciled with the interior composition information and machine-specific code is generated to drive a 3-D Printing machine. Demonstration parts with composition that varies from the surface toward the interior of the parts have been fabricated using this technology.

Introduction

This paper reports on a portion of our work under a project entitled, "Distributed Design and Fabrication of Metal Parts and Tooling by Three Dimensional Printing". The goal of this project is to identify and overcome barriers to the distributed design and fabrication of parts made by Solid Freeform Fabrication technology, using Three Dimensional Printing as a prototypical SFF process. One of the great potential benefits offered by Solid Freeform Fabrication technology is the ability to create parts that have composition variation within them. Such local composition control has the potential to create new classes of components. For example, monolithic components can be created, which integrate the function of multiple discreet components. Alternatively, material composition can be tailored within a component to achieve local control of properties (for example, increasing hardness, where needed, but leaving the material tough where the hardness is not required). The fine resolution of some SFF processes offers the opportunity for fabrication of a class of parts, which has been termed "mesoscopic" to connote parts that are smaller than typical mechanical parts, but not fabricated with microfabrication methods. Local composition control has a key role to play in such mesoscopic components, greatly increasing the potential of this application of the technology.

Among the SFF processes, Three Dimensional Printing is particularly well suited to the fabrication of parts with local composition control. 3-D Printing creates parts by spreading powder, and then ink-jet printing a "binder" material into the powder bed to define the component. In the past, our binders have often included materials to be printed, such as silica or metal particles in fine colloidal suspension. Three Dimensional Printing can be extended to the fabrication of LCC components by printing different materials in different locations, each through its own ink-jet nozzle or nozzles. Figure 1 illustrates this conceptually with three different colors, each representing the printing of a different material into the powder bed with local control of position. In this sense, Three Dimensional Printing of parts with local composition control is analogous to color ink-jet printing of documents, except that instead of controlling the color, you control composition, and instead of printing in 2-D, you print in 3-D. 3-D Printing is thus capable of fully three dimensional control of composition. Further, the resolution of composition control is on the order of the size of a primitive made from one droplet interacting with the powder bed—50-100μm. Such local composition control has been demonstrated in the printing of ceramic parts with local toughening [1]

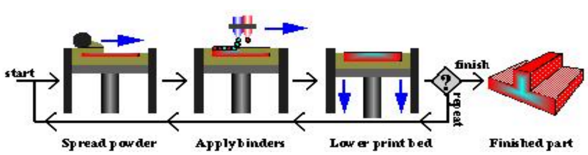

Figure 1. Local Composition Control with 3D Printing.

An attractive aspect of SFF technologies is that they can fabricate a very wide range of geometries due to the decomposition of 3-D geometry into 2-D layers and the subsequent fabrication as layers. This flexibility provides the opportunity to separate design from manufacturing, so that a design engineer can send a CAD model to a remote facility and have it successfully fabricated. Indeed, to a large extent, the rapid prototyping industry works in this manner. This is quite analogous to the separation of design and fabrication recognized by Carver Mead and Lynn Conway [2], which separation served as the engine for rapid growth of the integrated circuit industry. We seek to extend and preserve this "one-way information flow" to the domain of parts with local composition control. The advantage of one-way information flow is that successful products can be created without iteration and extensive discussions between designers and manufacturing engineers are not needed (thereby avoiding significant scheduling and logistical problems). Certainly, there will always be process limitations and constraints. In a manner analogous to the VLSI model, such constraints will be accommodated through design rules, wherever possible. These design rules capture these constraints in the simplest manner possible, so that the designer can operate within the limitations of the process.

This work develops the representations, methods and tools needed to achieve one-way information flow in the design of parts with local composition control. Designers need new CAD tools to capture their ideas as models with graded compositions and manufacturers need new data post-processing capabilities to convert these models into machine instructions for their fabrication. Design rules that capture the process limitation in a simple form that designers can use are also needed. The lack of these tools is the major technical barrier between the designers and an SFF-manufactured part with composition control.

This paper is organized as follows. The next section will present an overview of the information flow for the design and fabrication of a part with graded material composition. The third section will focus on the details of

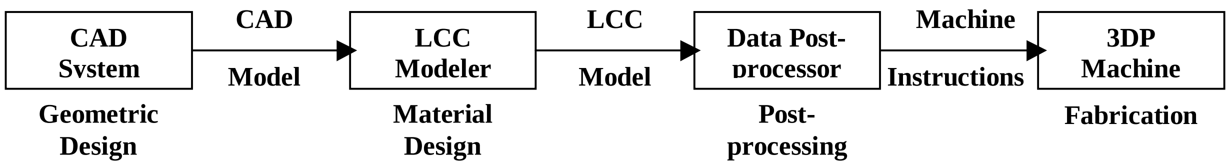

Figure 2. One-way Information Flow for LCC with 3D Printing

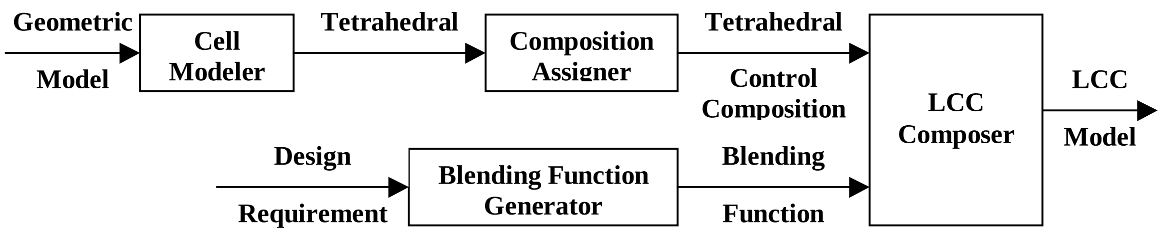

Figure 3. Data Flow of the LCC Modeler

the post-processing approach. Examples of the part with composition control are shown after.

Information Flow

The information flow for Local Composition Control with 3D Printing is shown in Figure 2. It is composed of four main steps: Geometric Design, Material Design, Post-processing and Fabrication.

The geometric design step generates the geometric data that define the physical boundary of the part. It is implemented through a commercial CAD system, such as SolidWorks.

The material design step is used to define material distribution within the physical boundary designed in the first step. For the design of a part with locally controlled composition, it is implemented through a module called an LCC Modeler. In the simplest approach, the designer will specify the composition of the material directly. In future approaches, the designer will be able to specify the desired properties (mechanical, electrical, etc.) and the translation to material composition will be implemented for the designer through relationships established in tabular form. This step finishes the design stage of the information flow and generates a model that reflects both the geometric and material design initials.

The Geometric data is exported from the CAD system as a standard interchange format such as IGES. The LCC modeler starts by creating a tetrahedral mesh (of the type used in FEM analysis). The internal composition will be established by prescribing the composition of the vertices of the tetrahedra and then interpolating between them. These vertices and their composition values are called control vertices and control compositions, respectively.

For an arbitrary point inside the geometric boundary of the part, its composition can be determined in the following manner. First, the bounding tetrahedron that contains this point is determined efficiently through a bucketing approach. The composition value of this point is then calculated by blending the control compositions associated with this tetrahedron through a predefined formula, e.g., the well-known barycentric Bernstein polynomials [3]. In this way, a continuous material system is defined in 3-D Euclidean space by the tetrahedral in a bucketing structure, the control composition for each vertex and the composition blending function. The definition is unique and integrated, as discussed in detail in [3,4] Various methods of efficiently specifying composition will be made available. At this time an efficient method is available for describing composition as it varies from the surface of a component.

The post-processing step is used to convert the LCC model generated by the design process into machine instructions that can be used to control the machine fabricating directly. The post-processing plays the role of interface between the design and fabrication stages. To fill this gap, the post-processing step needs to slice the LCC model to generate layer-wise information, transform the continuous material distribution into printable, discrete information through digital halftoning, merge the geometric information with the halftoned material information through printing patterns and generate machine instructions (see Figure 4). We will discuss the post-processing in detail in the next section.

The instructions produced by the last phases of the post-processing step are machine specific instructions, which direct the fabrication of the component. The machine executes on-line process control to maintain high accuracy in the execution of the post-processing

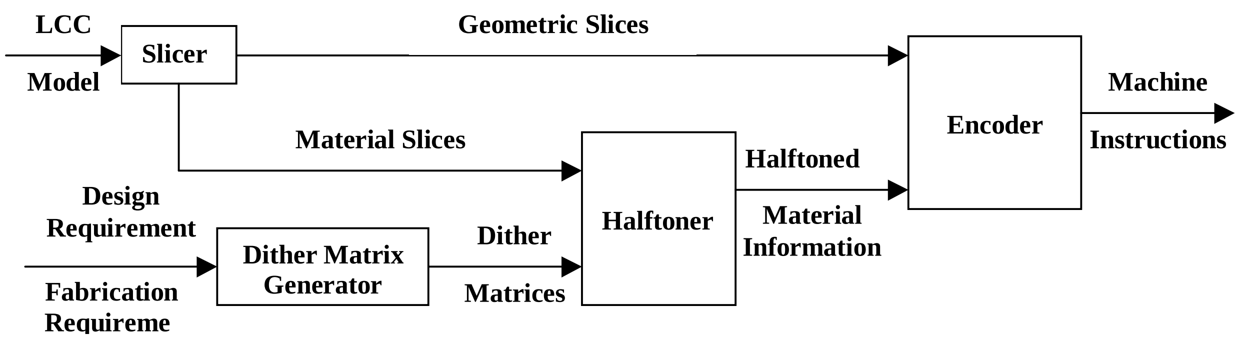

Figure 4. Data Flow for the Data Post-processing

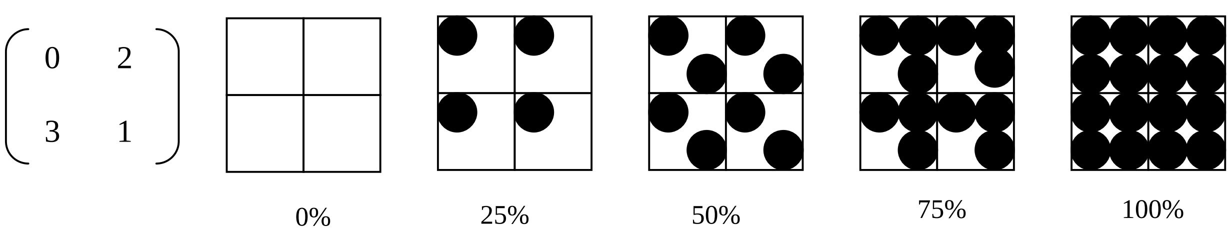

Figure 5. 2x2 Dither Matrix and the Corresponding Dither Patterns.

instructions. At the present time, the entire file is post-processed off-line and loaded onto the printing machine. This has the advantage of permitting the time consuming post-processing to be done off-line and the disadvantage of requiring significant file storage space. The alternative for future implementation may be on-the-fly post processing of a few layers at a time.

Data Post-Processing

In the data post-processing step, three major problems must be addressed, i.e. the halftoning, the merge of geometric and material information and the generation of machine instructions.

Digital Halftoning for LCC

As we mentioned above, the 3-D Printing process is analogous to ink-jet printing. It fabricates a part by adding binder into the powder in a drop-by-drop fashion (see Figure 2). Like ink-jet printing, 3-D Printing is a binary device. In other words, when a droplet is available, it is either printed or not printed. The material system is defined as continuous. To print a continuous material distribution on a binary device, a digital halftoning technique will is used.

Digital halftoning is a method of rendering the illusion of continuous-tone images on displays that are capable of producing only binary picture elements (PEL) [5]. Since most hardcopy devices render images as binary or discrete tones, halftoning is used to convert a continuous, ideal image within a computer representation into a form capable of being rendered on the hardcopy devices. A logical starting point for achieving LCC with 3-D Printing is to extend digital halftoning methods.

In the halftoning for color printing, dither matrices are generated to approximate the continuous-tone color scheme with different numbers of color dots at specific positions inside the area covered by a dither matrix. The size of a dither matrix determines how many levels of intensity it can generate. Let us consider gray scale only. For instance, the 2x2-dither matrix in Figure 5 can generate 5 different levels of intensity. If we consider all black a 100 percent of intensity, and all white a zero percent of intensity, then the intensity levels of zero, 25, 50, 75 and 100 percent of intensity can be generated by the dither patterns shown on the right of the dither matrix. Each of the intensities is illustrated by an area covered by four dither cells. A black dot represents a droplet. For more detailed information about the digital halftoning, please refer to [5]. The dither matrix on the left of Figure 5 determines the order of filling of a dither cell. When a 25% gray level is reached, the first droplet to turn on is the one in the upper left-hand corner ("0" in the dither matrix). When a 50% gray level is reached the droplet in the lower right hand corner ("1" in the dither matrix) turns on. At 75%, the upper right hand drop turns on. At full intensity, the lower left had droplet turns on.

The intensity level between any of the two adjacent levels can not be generated precisely. Normally a threshold scheme is defined to approximate these levels with the nearest precise levels. For instance, intensity level at 20 percent will be displayed as 25 percent.

In 3-D Printing with LCC, the dither matrix is used to generate different levels of material intensity. We only need to substitute the color intensity with material composition in the above concept and give different materials a different dither matrix. An algorithm for generating optimal dither matrices has been developed. This algorithm is based on Bayer’s algorithm [5], but has been adapted to meet certain requirements imposed by 3-D Printing.

Bayer's algorithm is based on the assumption that all PELs are in the same isotropic shape, which is quite reasonable for the traditional hardcopy printing in which these shapes are mainly circles, squares and hexagons. In 3-D Printing, a typical PEL is a 200μm high and 10μm wide rectangle, which is obviously too slim to be approximated by a isotropic shape.

What is more, due to the physical properties of the fluid as binder and the limits on the speed of the print head, there is a distance between two available binder droplets on the powder bed. If the frequency of the available droplets is 40kHz and the print head moves at 1.2m/s, for instance, the distance between two adjacent droplets is 30μm. In common cases, the droplet frequency and the print head speed will fluctuate around the pre-defined values with a small variation. It is hard to determine the precise location of the droplet inside the 30μm-wide region. On the other hand, the hardware resolution of the 3-D Printing machine is 10μm, i.e. the width of a PEL is 10μm. To make sure that there is always a droplet available when needed, we need to grant the three PELs in the 30μm-wide region identical value. This requirement imposed by the 3-D Printer is called the minimum run-length requirement.

The digital halftone algorithm for 3-D Printing starts from the same idea as Bayer's algorithm, while taking the above two requirements into consideration. It can optimally minimize the occurrence of low frequency textures and generatedither matrices for different minimum run-lengths and PELs with different height/width proportions under different printing parameters. Currently, the digital halftoning is imposed layer-wise, i.e. it is 2-D halftoning. To eliminate the low frequency effect in the Z direction, 3-D halftoning can be used and corresponding software is under development.

Intensity Levels	Patterns
0	0 0 0 0 0 0 0 0
1	1 0 0 0 0 0 0 0
4	1 0 0 0 1 0 0 0
8	1 0 1 0 1 0 0 0
14	1 0 1 0 1 0 1 0
16	1 0 1 1 1 0 1 0
20	1 0 1 1 1 0 1 1
26	1 0 1 1 1 1 1 1
28	1 1 1 1 1 1 1 1

Table I. Patterns for the First Row of the Dither Matrix in Figure 8.

Merge Geometric and Material Information

After the halftoning, the 3-D continuous-tone material definition has been transformed to 2-D discrete material composition levels. Now what we have in one layer is dither cells associated with composition values for different materials and geometric boundaries of the part (see Figure 6). One dither cell is the area covered by a dither matrix. Its composition value determines how to place material droplets in this area, as we discussed above. We need to combine the information describing the geometric boundary with the information describing the internal

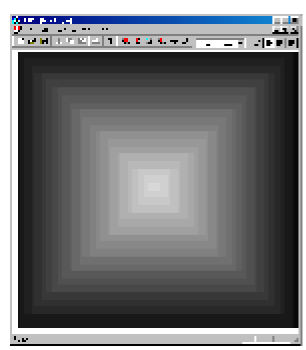

Figure 6. Dither Cells in One Layer of a Cube with LCC

composition. The image in Figure 6 is generated by using the composition value as gray scale value and displaying a gray rectangle for each dither cell. The composition is graded from 100 percent at the boundary to 15 percent at the center.

The first step is to allocate nozzles for each material. For instance, if we have two materials and eight nozzles, then we can allocate the first four nozzles to the first material and the remaining nozzles to the second material. When we are given a nozzle index, we can know which material it contains and use the corresponding dither cells.

Given a nozzle, a technique similar to the displaying of a filled polygon on a computer monitor is used. Since the electronic gun of a computer monitor draws only one line at a time, to display a filled polygon it needs to draw several raster lines. In one raster line, it starts lighting the pixels at the boundary of the polygon and stops lighting the pixels at the next boundary it meets and then iterates. In 3-D Printing, we also need to generate these raster segments for each nozzle. Rather than lighting the pixels, the print head prints a droplet inside the raster segment.

In LCC with 3-D Printing, since the composition varies along the raster line, we can not simply let the print head print. We need to assign the print head patterns according to the material composition inside the raster segment. The electronics on the 3-D Printing machine allow for the definition of repetitive patterns. The advantage of this approach is that these patterns need only be loaded once at the start of each pass of the print head and then can be called as needed during the pass. This is far more efficient than loading the information for each and every droplet. The 3-D printing machine has provisions for loading 8 different patterns for each nozzle. As will be discussed in the following paragraphs, this is sufficient to allow the specification of the 4 x 8 dither matrix of Figure 8.

To illustrate, Figure 7 shows a 2x2-dither matrix. There are altogether 4 different print head patterns: 00, 10, 01 and 11. Figure 8 shows a 4x8 dither matrix, together with the 11 patterns that are required to define all possible droplet patterns in an individual row. Fortunately, for a given row, only 8 patterns

$$\begin{pmatrix} 0 & 2 \\ 3 & 1 \end{pmatrix}$$

Intensity	0%	25%	50%	75%	100%
First Row	0 0	1 0	1 0	1 1	1 1
Second Row	0 0	0 0	0 1	0 1	1 1

Figure 7. Print Head Patterns of a 2x2 Dither Matrix in Each Row for Each Intensity

$$\begin{pmatrix} 0 & 28 & 8 & 16 & 4 & 26 & 14 & 20 \\ 6 & 24 & 12 & 22 & 2 & 30 & 10 & 18 \\ 1 & 29 & 9 & 17 & 5 & 27 & 15 & 21 \\ 7 & 25 & 13 & 23 & 3 & 31 & 11 & 19 \end{pmatrix}$$

Patterns:

10000000	10001010	10111011
00001000	10101010	11111011
10001000	10111010	11111111
10101000	10101011	

Figure 8. Different Patterns for a 4x8 Dither Matrix.

are needed. For example, table 1 shows the patterns needed to describe all possible combinations of the first row of the dither matrix of Figure 8. In general, a dither matrix that is *m* wide will require only *m* patterns to define a given row. As we mentioned above, once a dot is added into a dither cell in the halftone algorithm, it will never be removed or moved to another location. In one row of the $n\times m$-dither matrix, there are only *m* different positions. Every time a dot is added in this row, the pattern in the row is changed. Since we can only add *m* dots in an order determined by the dither matrix, we can only obtain $m+1$ different patterns in one row. One of the patterns is all 0's and this can be eliminated as it is always present. Thus, *m* patterns are sufficient.

The process of assigning proper print head patterns to raster segments is composed of two steps. The first step is to split a raster segment into several shorter segments so that in each new segment, there is only one pattern that will be used in the printing. This is a form of run length encoding and it greatly compresses the files produced. The second step is to generate the proper printing command for each segment according to the print head pattern associated to it.

Surface Finish in 3D Printing with LCC

The primary sources of surface roughness in 3D Printed parts is due to the discrete nature of the layers. In a raster machine, the lines that make up the layers are discrete as well. Fortunately, the ends of these lines can be blended into each other using the "proportional deflection" capability of the droplets on the 3-D Printing machine. The basic idea of proportional deflection is to offset the location of a droplet along the Y-axis. Since the raster line is printed along X-axis, this deflection function gives us the full control of the location of a droplet inside the 2-D area covered by the raster line.

Figure 9 illustrates the proportional deflection tracing of the edge. The solid lines are the boundary of the part. The overlap circles are binder droplets that will hold the powder together. The filled circles in the right figure are the droplets that have been deflected from their original positions, which are shown in the left figure. The thick curves show the result surface finish.

In 3-D Printing with Local Composition Control, we need to address a new kind of surface finish problem. As we discussed above, we are controlling the material composition by changing the number of droplets. If the composition intensity at the boundary of a part is too small, then we may

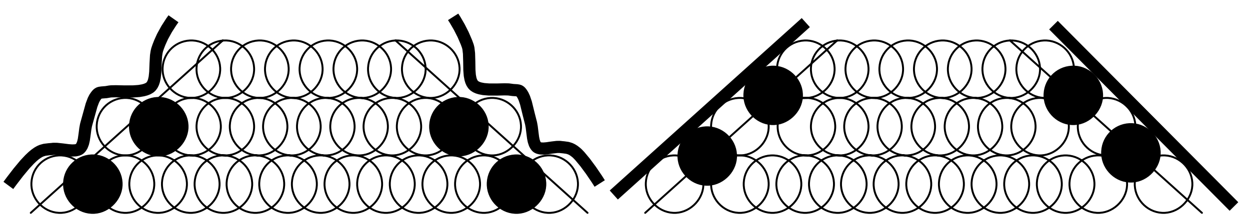

(a) Without Edge Tracing (b) With Edge Tracing

Figure 9. Edge Printing Without (a) and With (b) Proportional Deflection Edge Tracing.

not have enough droplets to control. For instance, if a has a composition intensity of 25 percent, then some of the dark circles may no

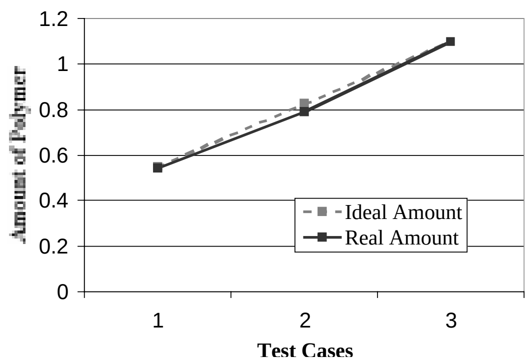

Figure 10. Comparison Between Theoretical and Test Data.

longer be available, according to the dither matrix shown in Figure 5. In this case, the deflection approach shown in Figure 9 must be modified.

Several approaches may be envisioned. We may force the binder concentration to be 100% at the boundary. For example, if a part is designed to have a composition profile graded in from the boundary it might start at 100% at the boundary and decrease moving into the part. Second, we can selectively add several extra droplets based on the real composition value. To do so, we need to have several criteria to determine how many droplets to add and in what manner to add them, so that we can get optimal surface finish while minimizing the violation to the original design. At the present time, we practice the first approach of enforcing 100% concentration at the boundary. The alternative will be developed.

Example Parts

The first group of test cases was used to check the precision with which composition can be rendered. Three bars of the same volume were printed with different composition values, i.e. 100 percent, 75 percent and 50 percent of the maximum allowable amount of polymer (defined by filling the pore space with printed material). After printing, the amount of polymer was measured by weighing the dried bars, burning out the polymer in a furnace and re-weighing the bars. As can be seen from Figure 10, the amount of polymer printed closely matches the instructions to the machine.

The second test case is used to test the whole information flow with a relatively general part. The shape of the part is shown in Figure 11. The material composition of this part is defined as starting from 100 percent of maximum at the boundary and grading to 15 percent at a location 5mm away from the boundary. The composition variation is linear with respect to the distance from the boundary. To illustrate the state of the inner area of the part, we only printed about one third of the total height. The result is shown in Figure 10, which shows the two identical parts printed side by side. The parts have not been taken out of the powder. From Figure 12 we can see that the color faded along with the composition and illustrates a clear composition variation as we expected.

Conclusions

Solid Freeform Fabrication technologies have the potential to create parts with local control of the composition. This work seeks to enable the distributed design and fabrication of such components, using Three Dimensional Printing as a prototypical SFF process. The "one-way" information flow from designer to fabrication process is described, which includes the steps of: (1) Geometry creation; (2) Design of the internal composition; and (3) Post-processing to produce code for the SFF machine and for fabrication by SFF. Previous work has described our efforts on the development of a CAD representation, which is capable of embodying local composition

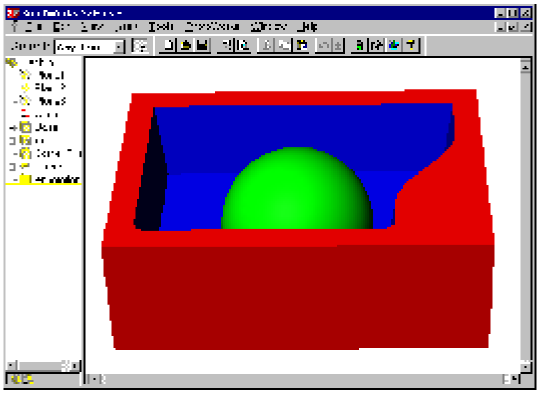

Figure 11. Shape of the Test Part

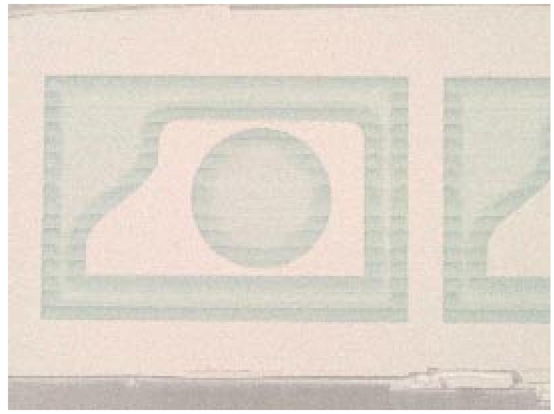

Figure 12. Printing Result

information. This paper focuses on the post-processing. The first step in post-processing is to render the desired continuous distribution of composition as a halftoned version, which can be created by the digital construction method of Three Dimensional Printing. Peculiarities of the Three Dimensional Printing machine, including anisotropic voxils and uncertainties in droplet placement are addressed through modifications to standard halftoning algorithms. In the next step, the halftoned information is compressed by run length and coating and efficiency is gained through the use of a standard number of "patterns". The use of these patterns substantially reduces the overhead at the 3-D Printing machine. A third consideration discussed involves reconciling the information about the geometric boundary and the interior composition where such conflicts occur. Finally, example parts printed by exercising the entire information flow are presented.

Acknowledgements

The authors gratefully acknowledge the National Science Foundation (#DMI-9617750) for their support of this research.

References

[1] J. Yoo, K. Cho, W. S. Bae, M. J. Cima, and S. Suresh, "Transformation-Toughened Ceramic Multilayers with Compositional Gradients", Journal of the American Ceramic Society, 81(1):21-32, January 1998.

[2] Carver Mead and Lynn Conway, Introduction to VLSI Systems. Addison-Wesley Publishing Company, Reading, MA, Oct. 1980.

[3] E. M. Sachs, N. M. Patrikalakis, M. J. Cima, D. Brancazio, W. Cho, T. R. Jackson, H. Liu, H. Wu and R. Resnick, "The Distributed Design and Fabrication of Metal Parts and Tooling by 3D Printing:", Proceedings of the NSF Design and Manufacturing Grantees Conference Proceedings, Long Beach, CA, January 1999.

[4] T. R. Jackson, N. M. Patrikalakis, E. M. Sachs, and M. J. Cima. "Modeling and Designing Components with Locally Controlled Composition", D. L. Bourell et al., eds., Proceedings of the Solid Freeform Fabrication Symposium, Austin, TX, August 1998, pp. 259-266.

[5] R. Ulichney. Digital Haltoning. MIT Press, Cambridge, MA 1987.