

Methods for Distributed Design and Fabrication of Parts with Local Composition Control

W. Cho¹ E. M. Sachs¹ N. M. Patrikalakis¹ M. J. Cima¹ T. R. Jackson¹
H. Liu¹ J. Serdy¹ C. C. Stratton¹ H. Wu¹ R. Resnick²

¹Massachusetts Institute of Technology, Cambridge, MA

²ExtrudeHone Corporation, Irwin, PA

Abstract

With the advances in Solid Freeform Fabrication (SFF), the ability to fabricate parts with Local Composition Control (LCC) is becoming a reality. Such LCC has the potential to create new classes of components. For example, monolithic components can be created, which integrate the function of multiple discrete components. Alternatively, material composition can be tailored within a component to achieve local control of properties (e.g., hardness vs. toughness, magnetic properties, corrosion properties, etc.). This article summarizes our recent work on the design of objects with LCC, their interrogation for visualization, and fabrication through post-processing and 3D Printing.

Introduction

One of the great potential benefits offered by *Solid Freeform Fabrication (SFF)* technology is the ability to create parts that have composition variation within them. Such *Local Composition Control (LCC)* has the potential to create new classes of components. Material composition can be tailored within a component to achieve local control of properties (e.g., index of refraction, electrical conductivity, formability, magnetic properties, corrosion resistance, hardness vs. toughness, etc.). By such local control, monolithic components can be created which integrate the function of multiple discrete components, saving part count, space and weight and enabling concepts that would be otherwise impractical. Controlling the spatial distribution of properties via composition will allow for control of the state of the entire component (e.g., the state of residual stress in a component). Integrated sensors and actuators can be envisioned which are enabled by LCC (e.g.,

bimetallic structures, in-situ thermocouples, etc.). Devices which have as their function the control of chemical reactions are possible. The utility of “Mesoscopic” parts made by SFF will depend strongly on the ability to locally control composition.

Realizing the potential utility of LCC in SFF is a many-faceted challenge requiring developments in the: (1) Information technology and design tools required to support the design of parts with LCC; (2) Extension and characterization of the range of materials which can be deposited with local control (SFF technology specific); (3) Design of materials systems with locally varying composition which can be successfully treated in operations subsequent to the SFF process itself (e.g., densified in a furnace firing operation); (4) Exploration of specific applications of LCC.

The work reported in this article focuses primarily on the issue of Information Technology and Design Tools – (1) above. The absence of knowledge, methods and tools in this area presents an absolute bar to the exploration of materials systems and applications. Developments in these areas will allow a wider community to contribute to materials and applications.

Information Technology and Design Tools may be divided into two categories: (1) tools which are generic, and (2) tools which are specific to a given SFF process. Generic electronic representations must be developed to allow for electronic specification within a component. There must be a suite of tools which allows a designer to communicate with this representation using high level features that are sensible to a designer. The designer must be able to visualize and interrogate the evolving model. The model must not allow the designer to request that which cannot be made. Process specific tools include methods to render desired continuous composition profiles in the discretized form required by a specific process and the generation of machine specific fabrication instructions.

Wherever possible, the work conducted under this

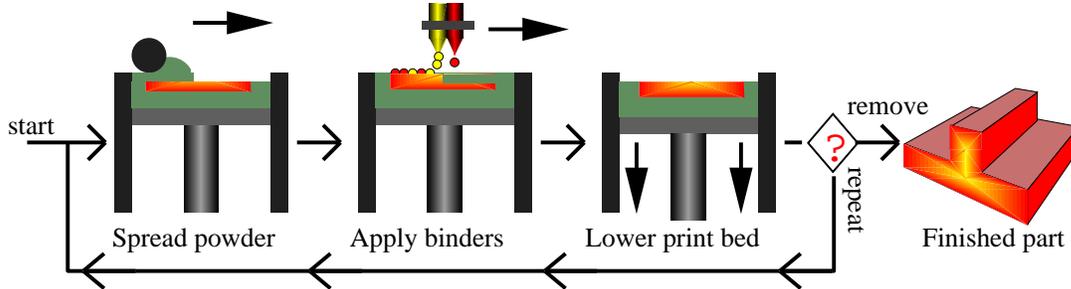


Figure 1: 3D Printing illustrating Local Composition Control (LCC)

project will be generic and applicable to a broad range of SFF technologies. However, in the cases where the outcome is process specific, *Three-Dimensional Printing (3D Printing)* will be used as the prototypical SFF technology. Among the SFF processes, 3D Printing is particularly well suited to the fabrication of parts with LCC. 3D Printing creates parts in layers by spreading powder, and then ink-jet printing materials into the powderbed [11, 12, 14, 15, 16]. In some cases, these materials are temporary or fugitive “glues”, but in many cases, these materials remain in the final component. Examples of the latter include; ceramic particles in colloidal or slurry form, metallic particles in slurry form, dissolved salts which are reduced to metal in the powderbed, polymers in colloidal or dissolved form, and drugs in colloidal or dissolved form. 3D Printing has been extended to the fabrication of LCC components by printing different materials in different locations, each through its own ink-jet nozzle(s). Figure 1 illustrates this conceptually with two different colors, each representing the printing of a different material into the powder bed with local control of position. 3D Printing is thus capable of fully three dimensional control of composition.

Information Flow

The current information flow for Local Composition Control (LCC) with 3D Printing is composed of four major steps: geometric design, material design, post-processing, and fabrication - see also Figure 2-(A). This information pathway establishes a clear separation between design of objects with LCC, their processing, and their fabrication. The geometric information is exported from the CAD system as a standard exchange format such as IGES [3] and STEP [2], followed by the generation of tetrahedral meshes. This process is referred to as *pre-processing* in Figure 2-(A). The internal composition will be established by specifying the composition values at the vertices of each tetrahedron and

interpolating between them. This finite-element based LCC modeler can be thought of as a special instance of our generalized cellular decomposition approach to LCC modeling [5, 6, 13]. It was chosen as a convenient method to demonstrate the information pathway and to explore the issues associated with LCC.

Post-processing then converts the designed LCC model into instructions for the 3D Printing machine. Post-processing takes place on a layer-by-layer basis along two parallel paths: (1) the accurate definition of the surface (geometric slice); and (2) rendering the composition of the body (material slice). The continuous-tone material composition is rendered into printable discrete information using a halftoning (or dithering) algorithm. The boundary and composition information is recombined to produce the drop-by-drop instructions that are loaded onto the 3D Printing machine. Special attention is given to reconciling conflicts which occur at boundaries where the designer’s intent in both composition and surface finish must be recognized.

The complete information and 3D Printing pathway has been tested and demonstrated with a part of representative complexity as shown in Figure 2-(B). The part is an injection molding tool and the design challenge is to place hard phases in a designed composition profile near the surface. In this demonstration, two colors of ink were printed (magenta and cyan) with the condition that the sum of the materials was everywhere constant. The bottom image in Figure 2-(B) shows a photograph of a layer of the actual printed part. This can be compared with the material and geometry information above it, which become merged to produce the instructions which led to the printed part.

Alternative Representations and Evaluations

We have analyzed various data structures for representing LCC objects including a voxel-based structure, finite-element mesh approach, and the generalized

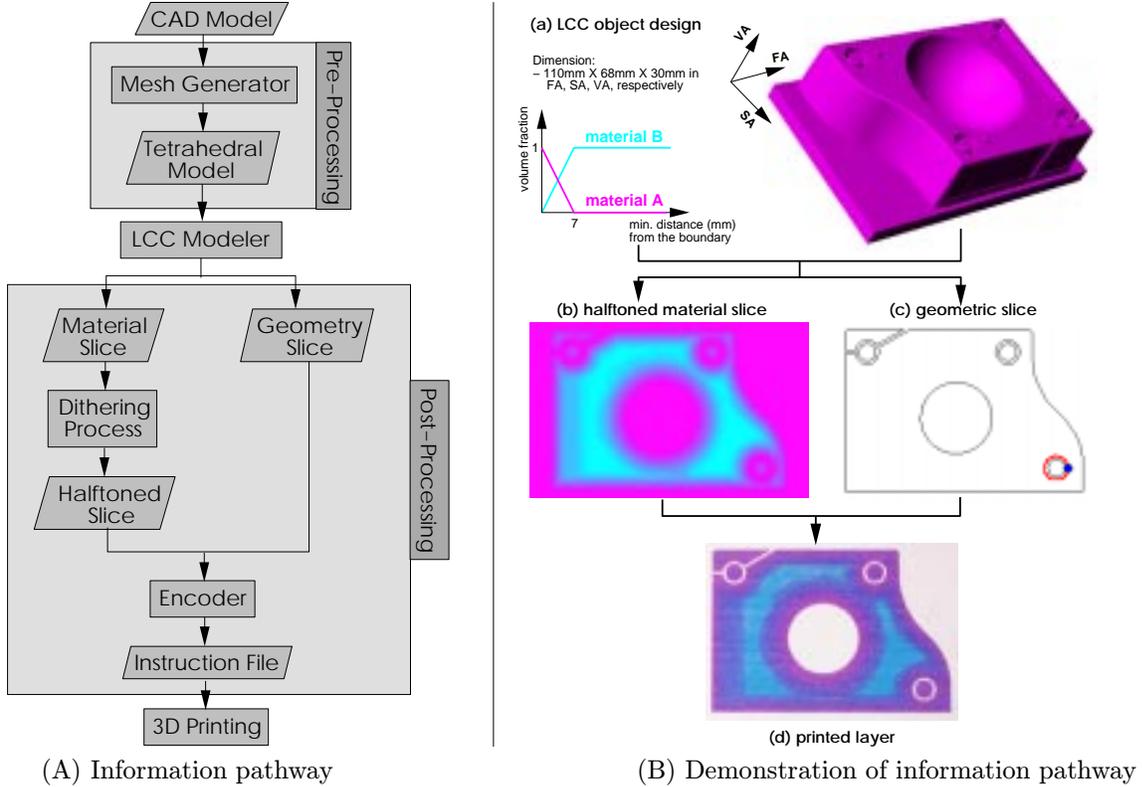


Figure 2: Information flow for LCC with 3D Printing

boundary representation (B-rep) data structures in order to represent spatially varying geometric/material properties [4]. A voxel-based modeling is an exhaustive enumeration approach, where a region of space containing the object is decomposed into a lattice of voxels and a vector of numerical values are associated with each voxel for blended compositions. In the finite-element mesh approach, material or physical property fields are attached to the nodes in the mesh. Interpolation functions (e.g., linear) associated with the volume elements (e.g., linear tetrahedra) are used to define the composition throughout each element as functions of the values assigned to the nodes. Generalized B-rep data structure maintains the topology of a model in a relational database and allows the incorporation of various geometric representations that best describe the geometry of the object’s model. This paradigm can be extended to the representation of LCC objects, where the geometry and composition is defined external to the topological data structure, allowing an accurate modular approach to the design of the LCC modeling system architecture.

The growth of storage cost in the voxelized method is sensitive to the size of the object, and the required geometric and material resolutions. The storage cost

of the finite-element mesh approach is a function of the prescribed geometric/material accuracies which depend on the minimum geometric feature size in a bounding surface, maximum geometric curvature of the bounding surface patches, and the minimum intended material feature size and the maximum material curvature within a region. The storage costs of the generalized B-rep data structures are constant with the desired accuracy of representation, and grow only with the number of features present in the model. Figure 4 illustrates a typical growth of the storage cost as the desired geometric/material accuracies vary. Figure 4-(a) shows the storage cost for each representation method with respect to geometric accuracy ϵ_g for material accuracy $\epsilon_m = 10^{-1}$. A material accuracy $\epsilon_m = 10^{-1}$ denotes the resolution of 10% in defining the composition. Similarly, Figure 4-(b) is the corresponding result for $\epsilon_m = 10^{-3}$. This analysis was performed for a realistic LCC object shown in Figure 3. Memory analysis of the tetrahedral mesh representation was performed for four minor variations of a typical tetrahedral mesh data structure, which results in a group of four graphs in Figure 4. Two kinds of boundary representations (cell-tuple-graph and radial edge) were considered for the analysis of storage cost for the generalized B-rep

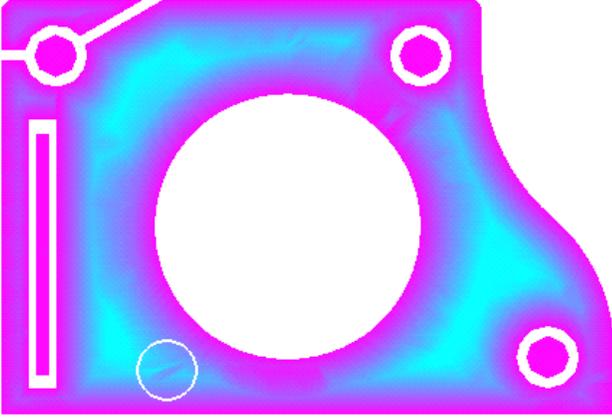


Figure 5: Material slice: linear composition variation with respect to the minimum distance from the boundary

- (almost identical solid and dashed lines in Figure 4). Each breakpoint (marked by a blue circle) in the graphs for the approximation methods, i.e., voxelized and tetrahedral mesh representations, shows the transition of the storage cost from being a function of geometric accuracy (ϵ_g) to being a function of material accuracy (ϵ_m).

In the course of this project we implemented an LCC modeler based on tetrahedral mesh data structure to provide a necessary link to design and process LCC models for fabrication. This was chosen as a convenient method to implement the information pathway and to explore the issues associated with LCC. The performance of this approach heavily depends on the quality of tetrahedral meshes. One major problem of the currently available tetrahedral mesh generators lies in the difficulty in controlling the local mesh size and shape. Tetrahedral meshing for complicated objects is a complex and memory intensive process, and it may require practically unaffordable processing cost to precisely describe the intended composition variation within a model. For example, an undesirable fingerprint (marked by a white circle) shown in Figure 5 is caused by the associated tetrahedron that is not fine enough and crosses the medial axis of the model. Similarly, the lack of composition variation within the three circular sections in Figure 5 also results from the coarseness of the meshing, i.e., no vertex near the center (degenerate medial axis) of the circle.

Another critical issue involved in the mesh-based approach is the lack of robustness in the tetrahedral mesh generation and the subsequent material slicing processes. This problem is unavoidable as exchanging complicated CAD models is error-prone, and errors are

magnified further in the subsequent complicated meshing/slicing operations mainly due to the imprecise computer arithmetic.

These drawbacks turned our attention to the voxelized representation scheme, which could serve as link between design and fabrication stages, and is algorithmically simple, robust, and facilitates the control of resolution. As a first step, we need to make sure that the storage cost of the voxelized representation falls within a practically affordable cost range for our parameter (ϵ_g and ϵ_m) values under consideration. We may want to set a voxel size equal to the size of a three-dimensional dither cell, which actually determines the geometric/material resolution of an LCC object. The size of a 3D dither cell is in fact arbitrary and its composition patterns can be determined by the corresponding 3D threshold matrices. Typically, a (4 by 2 by 2) 3D dither cell is appropriate in (fast, slow, and vertical) axis, respectively. The associated geometric resolution ϵ_g is $0.12mm$, $0.404mm$, and $0.34mm$ in each direction, if ($10\mu m \times 202\mu m \times 170\mu m$) PEL is used with a minimum run length of 3. For this (4 by 2 by 2) 3D dither cell, the material resolution ϵ_m is $1/(4 \times 2 \times 2) = 0.0625$. In other words, $\epsilon_g \in [0.12mm, 0.404mm]$, and $\epsilon_m = 0.0625$. Substitution of this (ϵ_g, ϵ_m) pair into our storage cost formula in [4] shows the storage cost for the voxelized representation is a magnification of the result in Figure 4-(a) by a constant scale factor 1.2. For $\epsilon_g \in [0.12mm, 0.404mm]$ (marked by a red line on the horizontal axis) in Figure 4-(a), the corresponding storage cost is approximately within the range of $[2 \times 10^6, 10^8]$ (also marked by a red line on the vertical axis). Therefore, rescaling this cost by 1.2 still guarantees the storage cost of the voxelized representation is within the order of 10^8 bytes. Efficient distance transform algorithms associated with this voxel-based modeling are also described in Section **Design Methods** below.

Pre-Processing

As an input to our system of design and interrogation of LCC objects, we consider a single solid represented by a boundary representation obtained from a CAD system, and exchanged via STL [1], IGES, or STEP file format. STL data is used for the extraction of geometric boundary information of an imported CAM model. IGES (or STEP) data is used as an input to tetrahedral mesh generation algorithms. Our data structure is based on the widely-used triangular finite-element meshing structure, and additional attributes such as a bucketing system and an association of facets with their parent subsets of the surface boundary re-

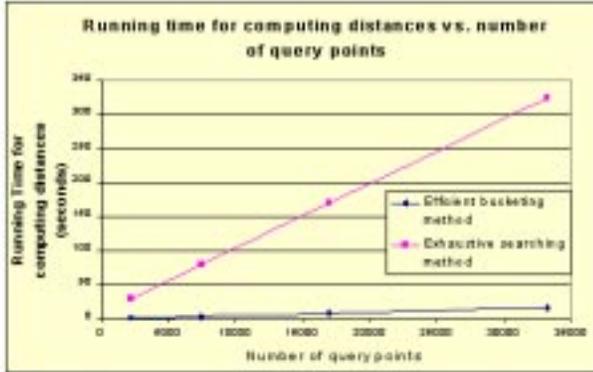


Figure 6: Comparison: time cost of distance computation for an LCC object in Figure 2-(B)

gion are included to enhance the efficiency in the evaluation and visualization of a composition.

Design Methods

We assume that the input geometry is a single solid represented via a boundary representation including tessellated models and curved models, obtained from a CAD system, and exchanged via a standard file format such as STL or STEP. The first algorithm we have already developed allows specification of the locally controlled composition as a piecewise polynomial or rational function of the minimum distance itself from the entire boundary surface. In order to design a composition as a function of minimum distance to the surface of the solid, an efficient distance transform is necessary. Among the various approaches to distance transform, the space division with rectangular lattice method is particularly useful and easy to implement. Specifically, the approach for improving efficiency of the distance transform includes preprocessing the model with bucket sorting and a digital distance transform of the buckets [8, 9]. The specific digital distance metric used was chessboard digital distance. Complexity analysis of the algorithm outlined above and experimental results demonstrate effective performance as shown in Figure 6.

Where the need for exact evaluation of the Euclidean distance from points within a solid to its boundary can be relaxed, an approximation based on Euclidean digital distance is useful and was investigated. With the model discretized into a large number of subdivisions, the Euclidean digital distance is potentially a very good approximation of the exact distance. Euclidean digital distance transform (EDT) algorithm developed by Saito & Toriwaki [17], which utilizes certain spatial coherence

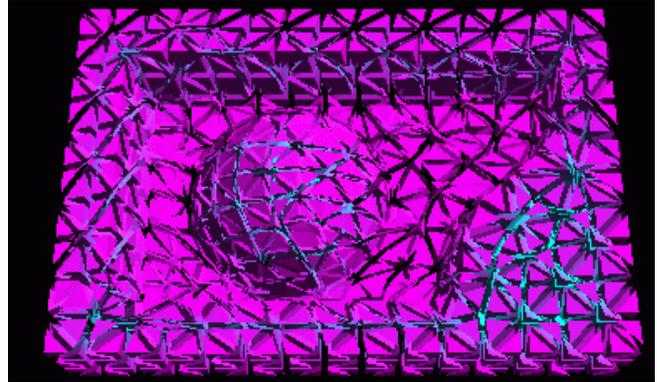


Figure 7: Visualization of a cuberille method for an LCC object in Figure 3

in scanning direction to improve efficiency, is one of the fastest EDT algorithms.

We also plan to develop an algorithm for the specification of FGM composition as a piecewise polynomial or rational function of the minimum distance from the boundary and also varying according to the boundary facet from which the minimum distance is computed from. It is well-known that the distance function from the boundary of a 3D solid is continuous everywhere. It is also differentiable except on the medial surfaces. Once the designer specifies the composition functions for each material (in terms of the minimum distance and also the footpoint), then the problem reduces to developing algorithms for efficient evaluation of composition at either arbitrary points within the solid or a sequence of points that exhibit spatial coherence. The above design methods have possible applications in new turbine blade and heat exchanger designs.

Visualization

Once the composition function is evaluated effectively, the visualization of the composition is done through various standard computer graphics techniques. The methods implemented include color-coded point sets, color-coded planar sections, cuberilles, and ray casting of the composition [9, 4] - (see also Figures 5 and 7 for the visualization of a cross section, and a cuberille method, respectively.)

Volume Dithering

In the information flow for local composition control (LCC) with 3D Printing, a dithering algorithm plays an important role as it converts an idealized continuous-tone LCC representation into a discrete version of ma-

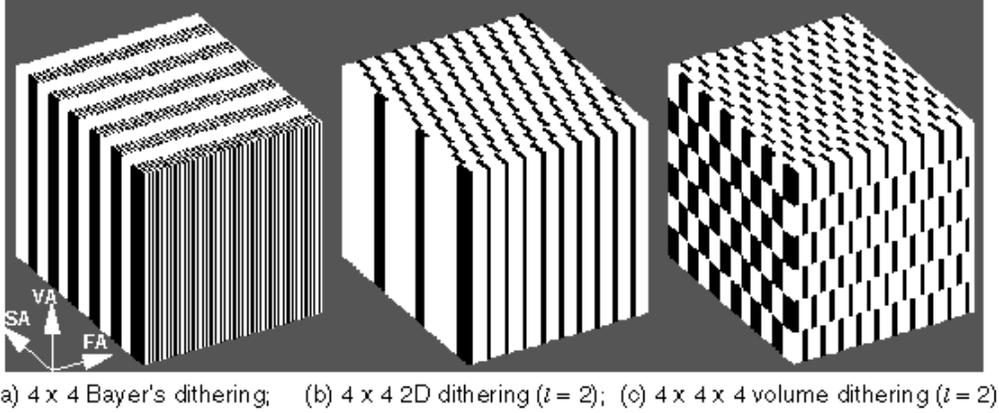


Figure 8: Ditherings within a volume ($2mm \times 2.4mm \times 2.1mm$) of uniform composition (25% black & 75% white)

chine instructions. Inspired by the fact that 3D Printing is analogous to ink-jet printing, our dithering algorithm is based on the classical digital halftoning technique [18] and adapted to minimize the undesirable low frequency textures of composition not only layer-wise but throughout the volume of an LCC model - see also Figure 8-(c) for the comparison with other dithering approaches in Figures 8-(a,b).

Compositions over the adjacent layers are sampled and compared with 3D threshold matrices. A binary representation approximating the original, continuous-tone values is output and then translated into machine instructions for fabrication. Due to the physical properties of the fluid as binder and the limits on the speed of the printhead, there is a distance between two available binder droplets on the powder bed. For instance, if the frequency of the available droplets is 40KHz and the printhead moves at $1.2m/s$, the distance between two adjacent droplets is $30\mu m$. As the droplet frequency and the printhead 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 $30\mu m$ -wide region. Considering the hardware resolution (i.e., the width of a picture element, PEL) of the 3D printer is $10\mu m$, we need to assign the three PELs in the $30\mu m$ -wide region an identical value. This is called minimum run length (l) requirement, and handled in the dithering process by using a dither PEL of width ($l \times$ hardware resolution) or $30\mu m$.

The volume dithering algorithm also takes into account certain technical limitations in the machine, only generating lattices that can be represented within the memory limits of the current hardware implementation in the 3D Printer. For example, the current 3D Printing machine has provisions for loading eight different patterns for each nozzle. Considering the feature

of dithering patterns¹, the limitation in the available machine patterns restricts the number of dither PELs to eight, only along the fast axis of printhead motion, while we can still achieve as high a material resolution as possible by controlling the size of dither cell in the other directions.

Encoding

The encoding process is composed of three main steps: rastering the geometric slice, generating a state map of the pixels for each material, and encoding the state map into printing instructions.

The raster of a geometric slice is an analogy to the technique used in displaying a filled polygon on a CRT screen. Since the electron gun of a CRT 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, then iterates. In 3D Printing, we also need to generate these raster segments for each nozzle. Rather than lighting pixels, the nozzle prints droplets inside the raster segment.

With the raster segments generated, we then match each segment with the halftone result of the material. This will further break the raster segment into several shorter segments, known as Materialized Raster Segment (MRS). Each MRS covers the consecutive dither-cell rows with identical Dither-cell-row Pattern. Each dither cell is composed of several rows. Each row has a determined state or pattern as of the states of all the

¹Once a dot is added into a dither cell, it will never be removed or moved to another location as the material density increases [19].

PELs in this row (filled or not) for a given composition. This pattern is called a Dither-cell-row Pattern. Two adjacent dither cells in the halftone result may have different composition values, but they may have the same Dither-cell-row Pattern for a certain row. In this case, that row in the two dither cells will be put into the same MRS. After processing all raster segments on one material, we actually obtain a run-length representation of a state map of all the pixels for this material in a layer. As a result, we now have a state map for each material on each layer.

Next, we encode the MRS's into machine instructions based on the pattern capability of the 3D Printing machine, (the α -machine at MIT). It is quite straightforward to transform a Dither-cell-row Pattern into a nozzle pattern. At first the materials must be allocated to different nozzles. Due to the design of the printhead on the α -machine, the nozzles have to be allocated in a pair-wise manner. If we have four materials, for instance, we would allocate nozzles 0 and 1 to material A, nozzles 2 to 3 to material B, etc. If we only have three materials, then nozzles 6 and 7 will be idle. No printing instruction will be generated for these nozzles. Since there are only 8 nozzles on α -machine, this imposes another limitation on the number of material we can have. For α -machine, at most four materials can be printed at the same time.

After allocating nozzles to material, a map from the nozzle index to the state map that should be used when generating the instructions is established. Then the software simply determines for each given printing time which part of the state map each nozzle should use, and generates proper instructions from it.

At this stage, an instruction file is generated and output in a binary format that can be interpreted by the controlling software on the α -machine.

3D Printing

3D Printing of parts with local composition control is accomplished using the MIT α -machine. The primary characteristic of this machine is that it has an 8-jet continuous-jet printhead, which defines the part by raster scanning over the powderbed. Continuous-jet printing is the preferred technology for accomplishing local composition control, because it allows for the widest range of options in binder materials to be printed. Continuous-jet printing has demonstrated the capacity to print aqueous and solvent-based materials, colloids and slurries, and dissolved matter. The one negative factor associated with the use of continuous-jet printing for local composition control is that with the current technology, there is an uncertainty of 1

droplet in any run length. In other words, if we ask for 2 sequential droplets of material A, we may get as many as 3 or as few as 1. This factor has been accommodated by developing dithering algorithms which allow us to specify minimum run lengths (see Section *Volume Dithering* above). It should be noted that in desktop ink-jet printers, color printing is accomplished with drop-on-demand type printheads. While these printheads do allow for the accurate specification of run length, including printing single droplets, they do not allow for the use of a wide range of materials.

The 8-jet printhead on the α -machine was configured with two banks of 4 jets for each material. Special manifolds and fluid recirculation components were fabricated to keep the two fluids isolated. For demonstration purposes (see also Figure 2-(B)), printing was accomplished using colored inks, cyan and magenta, on a white alumina ceramic powder, so that the colors would show up well. In this case, the binder carried color only, and a dissolvable, polymeric powder (polyvinyl alcohol: PVA) was added to the powder. When the aqueous dye hits the powder, it locally dissolves some of the PVA in order to bind the part, while the printed dye imparts color.

The α -machine has the capability to store the definition of 8 patterns for each of the jets. As explained above, the halftoning algorithms are designed so that in any given pass of the printhead, only 8 patterns will be required. However, the full definition of the image requires more than 8 patterns, and so at the beginning of each pass, some of the patterns may have to be overwritten in the pattern memory for each nozzle. At the present time, this pattern writing slows the process down considerably, resulting in a total time to print each layer of approximately 5 minutes (contrast this with less than 1 minute for single material printing). Future systems with more pattern memory can, of course, greatly increase printing speed.

Conclusions

Most CAD research has focused on the representation of 3D geometry, on methods and tools for designers to interact with these representations at a high level, and on derivation of machine specific instructions for machining. This work proposes the creation of complementary capabilities in the area of LCC and SFF. It is hoped that the availability of such tools and methods will be the key ingredient needed for the community to explore the potential of LCC.

This article has described our recent work on *Distributed Design and Fabrication (DDF)* via 3D Printing which enables a designer to create a design, send it

electronically for fabrication and have a part returned which conforms to expectations *without the need for iteration*. Our work has sought to identify and overcome the barriers to such DDF and to demonstrate accomplishment with parts made by 3D Printing. The major barriers identified include the inability of designers to exploit LCC due to the lack of electronic representations and design tools.

At this time, it is difficult to predict the eventual impact of LCC on the practice of engineering. SFF processes must become broadly used for *manufacturing and not just prototyping* (an active area of research within the 3D Printing project). Materials systems which can be processed to produce finished components with locally varying composition without distortion must be developed. Compelling applications must be demonstrated. These are significant risks.

Encouragement can be taken from the fact that in the 3D Printing project, several promising applications are under active development. Drug delivery devices are being created by printing different drugs at prescribed locations within the interior of a pill or implantable device. These drugs are then released into the body according to designed release profiles [10, 7]. A new program has just begun on Gradient Index Lenses (**GRIN**) which refract light by gradients in the index of refraction, rather than by external geometry. Such lenses can provide the functionality normally associated with multi-component ground optics at lower cost and in a smaller space. The drug delivery and GRIN applications are for high value added devices which are small in size and thus can reasonably be manufactured by 3D Printing. LCC is also being applied to the fabrication of tooling by 3D Printing. Hard phases such as TiC are being printed local to the surface of a tool for increased wear resistance. Tools with local control of porosity (for venting of gasses) are being fabricated by printing a material which acts to block the infiltrant during furnace densification. Although large in size, tooling applications can be economical because small quantities are required.

While one may not hope to match the impact of VLSI fabrication methods on engineering and society, the parallels are intriguing. VLSI and SFF are layer processes. VLSI depends on local control of composition and SFF is capable of the same. Perhaps, as in the case of VLSI, we will find that designers, given the proper tools, will find uses not now imagined for LCC in SFF.

Acknowledgment

Funding for this work was obtained from NSF and ONR under grants DMI-9617750 and N00014-00-1-

0169.

References

- [1] 3D Systems, Inc. *Stereolithography Interface Specification*, July 1988.
- [2] American National Standards Institute. *Product Data Exchange Using STEP (PDES) Part 42, Integrated generic resources: geometric and topological representation*. Fairfax, VA, February 1995.
- [3] IGES/PDES Organization, U.S. Product Data Association, Fairfax, VA. *Digital Representation for Communication of Product Definition Data, US PRO/IPO-100, Initial Graphics Exchange Specification (IGES) 5.2*, November 1993.
- [4] T. R. Jackson. *Analysis of Functionally Graded Material Object Representation Methods*, PhD thesis, Massachusetts Institute of Technology, January 2000.
- [5] T. R. Jackson, H. Liu, N. M. Patrikalakis, E. M. Sachs, and M. J. Cima. Modeling and designing functionally graded material components for fabrication with local composition control. *Materials and Design*, 20(2/3):63–75, June 1999.
- [6] T. R. Jackson, N. M. Patrikalakis, E. M. Sachs, and M. J. Cima. Modeling and designing components with locally controlled composition. In D. L. Bourell et al, editor, *Solid Freeform Fabrication Symposium*, pages 259–266, Austin, Texas, August 10-12 1998. The University of Texas.
- [7] W. E. Katstra, R. D. Palazzolo, C. W. Rowe, B. Giritlioglu, P. Teung, and M. J. Cima. Oral dosage forms fabricated by Three-Dimensional Printing™. *Journal of Controlled Release*, 66(1):1–9, May 2000.
- [8] H. Liu, W. Cho, T. R. Jackson, N. M. Patrikalakis, and E. M. Sachs. Algorithms for design and interrogation of functionally gradient material objects, *Proceedings of 2000 ASME DETC/CIE, 26-th ASME Design Automation Conference*, September, 2000, Baltimore, Maryland, USA. p.141 and CDROM, NY:ASME, 2000.
- [9] H. Liu. *Algorithms for Design and Interrogation of Functionally Graded Material Solids*, Master's Thesis, Massachusetts Institute of Technology, February 2000.

- [10] C. W. Rowe, W. E. Katstra, R. D. Palazzolo, B. Giritlioglu, P. Teung, and M. J. Cima. Multi-mechanism oral dosage forms fabricated by Three-Dimensional PrintingTM, *Journal of Controlled Release*, 66(1):11–17, May 2000.
- [11] E. M. Sachs, E. Wylonis, S. Allen, M. J. Cima, and H. Guo. Production of injection molding tooling with conformal cooling channels using the three dimensional printing process, *Polymer Engineering Science*, 40(5):1232–1247, May 2000.
- [12] E. M. Sachs, D. Brancazio, J. Milner, J. Serdy, A. Curodeau, and J. Bredt. High rate, high quality 3D printing through machine design, on-line measurement, and control, *International Journal of Machine Tools and Manufacturing*, to appear.
- [13] E. M. Sachs, N. M. Patrikalakis, D. Boning, M. J. Cima, T. R. Jackson, and R. Resnick. The distributed design and fabrication of metal parts and tooling by 3D Printing. In *Proceedings of the 1998 NSF Design and Manufacturing Grantees Conference, Cintermex Conference Center, Monterrey, Mexico*, pages 35–36. Arlington, VA: NSF, January 1998.
- [14] E. M. Sachs, J. Haggerty, M. J. Cima, and P. Williams. Three-Dimensional Printing Techniques, U.S. Patent No. 5204055, April 20, 1993.
- [15] E. M. Sachs, M. J. Cima, P. Williams, D. Brancazio, and J. Cornie. Three dimensional printing: rapid tooling and prototypes directly from a CAD model, *Journal of Engineering for Industry*, 114(4):481-488, November 1992.
- [16] E. M. Sachs, M. J. Cima, J. Bredt, and A. Curodeau. CAD-casting: direct fabrication of ceramic shells and cores by three dimensional printing, *Manufacturing Review*, 5(2):118-126, June 1992.
- [17] T. Saito and J. Toriwaki. New algorithms for Euclidean distance transformation of an n-dimensional digitized picture with applications, *Pattern Recognition*, 27(11):1551–1565, 1994.
- [18] R. Ulichney. *Digital Halftoning*, MIT Press, Cambridge, MA, 1987.
- [19] H. Wu, E. M. Sachs, N. M. Patrikalakis, D. Brancazio, J. Serdy, T. R. Jackson, W. Cho, H. Liu, M. J. Cima, and R. Resnick. Distributed design and fabrication of parts with local composition control. *Proceedings of the 2000 NSF Design and Manufacturing Grantees Conference*, Vancouver, BC, Canada, January 2000.