3D Hatching: Linear Halftoning for Dual Extrusion Fused Deposition Modeling

Tim Kuipers
Ultimaker
Watermolenweg 2
Geldermalsen, Gelderland 4191 PN
Netherlands

Eugeni Doubrovski
Delft University of Technology
Delft, Zuid-Holland 2628 CE
Netherlands

Jouke Verlinden
Delft University of Technology
Delft, Zuid-Holland 2628 CE
Netherlands

Figure 1: 3D prints obtained by applying vertical and diagonal hatching on a 14 cm 3D scan, a 15 cm artistic figurine, a full size soda can with textual information, and the result of a stress analysis performed on a 16 cm connecting rod of a piston engine.

ABSTRACT

This work presents halftoning techniques to manufacture 3D objects with the appearance of full grayscale imagery for Fused Deposition Modeling (FDM) printers. While droplet-based dithering is a common halftoning technique, this is not applicable to FDM printing, since FDM builds up objects by extruding material in semi-continuous paths. A set of three methods is presented which apply a linear halftoning principle called 'hatching' to horizontal, vertical and diagonal surfaces. These methods are better suited to FDM compared to other halftoning methods: their applicability stands irrespective of the geometry and surface slope and the perceived tone is less sensitive to the viewing angle. Furthermore, the methods have little effect on printing time. Experiments on a dual-nozzle FDM printer show promising results. Future work is required to optimize the interaction between the presented methods.

CCS CONCEPTS

•Applied computing → Computer-aided manufacturing;

KEYWORDS

FDM, 3D printing, dual extrusion, color, tone, grayscale, monochrome, linear, halftone, dithering, hatching, engraving

ACM Reference format:

Tim Kuipers, Eugeni Doubrovski, and Jouke Verlinden. 2017. 3D Hatching: Linear Halftoning for Dual Extrusion Fused Deposition Modeling. In *Proceedings of SCF '17, Cambridge, MA, USA, June 12-13, 2017*, 7 pages. DOI: <http://dx.doi.org/10.1145/3083157.3083163>

1 INTRODUCTION

This paper presents a process for fabricating 3D grayscale objects using the Fused Deposition Modeling (FDM) 3D printing method. It uses a principle which is based on modulating the visible width of printed lines of two alternating colors to produce the appearance of continuous tone gradients. The principle of creating the perception of continuous tones by printing small patterns of discrete colors is termed 'halftoning'.

The ability to apply color to 3D printed parts is relevant for both prototyping and manufacturing. Possible applications include reproduction of color-scanned 3D objects and fabrication of products with logos and labeling. Color can also be used as a design feature or to visualize geometric information such as the results of finite element analyses. See figure 1.

At present, 3D printing in color is available for a variety of Additive Manufacturing (AM) systems that are predominantly based on ink-jet technology. However, printing objects with high frequency color details using the FDM principle itself, without additional hardware such as ink-jet heads, was not available.

This work is licensed under a Creative Commons Attribution International 4.0 License.

SCF '17, Cambridge, MA, USA

© 2017 Copyright held by the owner/author(s). 978-1-4503-4999-4/17/06...\$15.00
DOI: <http://dx.doi.org/10.1145/3083157.3083163>

Implementing color variation with high frequency details using FDM is a challenge. FDM builds up objects by extruding material in semi-continuous paths, which makes it impossible to apply droplet-based halftoning principles that are commonplace in existing color 3D printers.

A promising approach to fabricate continuous tone objects using FDM has been presented by Reiner et al.. However, since that technique inherently produces textures at a relatively low sample rate, it does not allow the fabrication of high frequency details. Furthermore, the approach does not allow fabrication of textures on horizontal surfaces and degrades for diagonal surfaces with a slope approaching horizontal.

Addressing these issues, we propose a novel approach for 3D halftoning for dual-extrusion FDM that combines distinct strategies for horizontal, vertical, and diagonal surfaces. The proposed halftoning principle is based on hatching, an established 2D halftoning principle based on lines instead of dots. The implementation of the methods described in this paper is open source and can be found at github.com/Ultimaker/CuraEngine [Kuipers 2017].

2 RELATED WORK

2.1 Commercial color 3D printers

The first commercial full color 3D printing systems date back to 1993 [Corporation 2005]. These systems use ink-jet technology to apply colored binder onto white powder [Iliescu et al. 2009]. Consecutive layers of bound powder form the final 3D model. Instead of jetting a binder onto a substrate, Mcor developed a process in which conventional ink is jetted onto sheets of paper, which are then cut and stacked [Mcor 2013]. Stratasys uses ink-jet technology to print the building material itself. Their recent system incorporates six heads, each able to print a colored material [Stratasys 2016]. More recently, HP Inc. introduced a printing technology in which liquid agents are jetted onto powder in order to alter the powder's fusing behavior. According to the company, these agents may also include color in the future [HP 2014].

2.2 3D Halftoning

Because printers work with a limited number of base colors, specific strategies need to be applied in order to make full color prints. In 2D printing, this is usually done through different halftoning techniques.

2D ink-jet technologies apply a halftoning principle called dithering. In dithering, the distance between printed colored dots is varied to create perceived variations of colors. While halftoning for 2D printing industry is well developed, halftoning for 3D printers is still an active field of research. The first mention of halftoning in 3D printing is not focused on color reproduction but on material density variation for stereolithography [Lou and Stucki 1998]. Principles for 3D color dithering have been presented for binder jetting printers [Cho et al. 2003]. Vidim et al. presented a programmable pipeline for multi-material 3D printing. In their pipeline, dithering is applied for both visual and mechanical properties [Vidim et al. 2013]. With the aim to produce full-color prints using material jetting technology, Brunton et al. presented a halftone approach that takes into account the translucency of the printed material.

The different commercial systems discussed above all use ink-jet technology, and the color halftoning methods considered are based on the ability of ink-jet to deposit discrete droplets of color in the micrometer range. This allows the production of high frequency details. However, FDM lacks the ability to deposit discrete features in this size range, since it builds up objects by extruding semi-continuous lines of material. Therefore, the halftoning principles discussed above cannot be directly adapted to FDM. To create high frequency details using FDM, the development of different halftoning approaches is needed.

2.3 Color Fused Deposition Modeling

Limited by the semi-continuous material extrusion principle of FDM, approaches have been presented that aim to reproduce continuous tones in FDM prints. These can be categorized into two tactics. (1) Mixing material and color prior to extrusion. (2) Applying halftone principles.

Corbett's FDM grayscale printing technique performs continuous color mixing by implementing a mixing nozzle in the print head [Corbett 2012]. In this setup, multiple feeders are connected to a single nozzle. Multiple materials are fed into a volume of the nozzle where they are molten and mixed. The main limitation of this approach is that abrupt color changes require the full volume inside the mixing nozzle to be flushed.

Reiner et al. have shown that a modified dithering principle can be applied to dual-nozzle FDM systems in order to produce two-tone texture mapped 3D prints [Reiner et al. 2014]. Their technique involves applying sine patterns to the outer contour of each layer, with each layer alternating amplitude and alternating between black and white filament. A texture-based amplitude modulation is then applied in order to make the dots of the one filament protrude more than the other, which results in a shift in perceived color toward the former filament. The main challenge of this technique is to align the phase of the sine pattern across different layers that have a different geometry. The horizontal resolution of the produced textures is limited by the sine's wavelength, which in turn is limited by the width of the extruded lines. The application of the sine pattern therefore results in loss of high frequency details in the geometry. Moreover, this approach works best for vertical surfaces and performs significantly less on near-horizontal surfaces.

3 METHOD

Hatching is a halftoning principle which dates back to 17th century engraving techniques [Lavin 2004]. An example of this can be seen in figure 2a. Variations in the perceived tone are achieved by varying the local ratio between the width of black lines and the width of the white background surrounding it [Praun et al. 2001]. In more recent developments, image processing algorithms have been proposed which convert grayscale images into black and white engraving-style halftone images [Yamamoto 2006]. Their results resemble figure 2b, which was produced using the linear Newsprint filter from the GNU Image Manipulation Program. More recently, Freudenberg et al. adapted a hatching technique for the artistic rendering of 3D meshes [Freudenberg et al. 2002]. The linear characteristic of hatching makes it a particularly suitable approach for halftoning in FDM.

(a) Closeup of 'The Sudarium' [Mellan 1649]. ^a^aThis image is in the public domain.

(b) Engraving-style halftone image.

Figure 2: Examples of 2D hatching.

This paper presents a hatching approach for 3D models where three distinct techniques are applied depending on the slope of the 3D surface. Hatching techniques for horizontal, vertical and diagonal surfaces differ in the method by which the visible width of the extruded lines is modulated. Different strategies are needed since the visibility of printed paths varies strongly under different slopes of the surface. Because in FDM the objects are built up in discrete planar layers, a surface exposes either the top of a layer, the sides of layers, or both. Each of the techniques are discussed separately in the following sections.

The overall patterns by which conventional slicing software applications generate semi-continuous lines to build up a 3D printed object are unaltered by the proposed hatching techniques. Each layer contains *walls*, which are the paths following the outer contours of the object. The areas within those walls which are near the shell of the 3D mesh are filled by dense top/bottom *skin*. See figure 3. The layers are printed with alternating black and white filament and the amount of visible area of one layer with respect to the adjacent layers is altered in order to vary gray scale values.

3.1 Horizontal Hatching

Horizontal top and bottom surfaces of 3D meshes are hatched by modulating the widths of the lines used to print the top/bottom skin.

Figure 4: Horizontal hatching of a linear gradient. The 'c' marks where the cross section of figure c is located. The dashed lines in b show the travel paths of the center of the nozzle while printing the black lines. c shows the model used for achieving lines of a given width.

Figure 3: Semi-continuous lines in FDM.

The most commonly used pattern to fill in these areas is a regular grid of equidistant straight lines. Gray scale gradients are achieved by modulating the width of these lines. When printing a black layer the line widths are modulated so that parts of the adjacent white layer become visible and vice versa. A visual representation of this is presented in figure 4b. The resulting halftoning images resemble the image in figure 2b.

The texture image is sampled at regular intervals along the skin lines and the width w of each line segment between two consecutive sample points is determined by the line distance d and the average texture lightness L , which is in the range $(0, 1)$: $w = Ld$.

Varying the width of lines is achieved by increasing and decreasing the amount of deposited material. For Bowden style FDM printers it takes relatively long to change the amount of material departing the nozzle per second - a.k.a. the *flow*. Therefore we propose keeping the flow constant and varying the movement speed v of the print head: $v = c/A$, where A is the area of a cross section of the printed line and c is an empirically determined constant flow. We model the cross section as a rectangle with semicircular sides: $A = \pi(\frac{1}{2}h)^2 + h(w - h)$, where h is the layer height. See figure 4c. For lines narrower than the layer height we model the cross section as a circle: if $w < h$ then $A = \pi(\frac{1}{2}w)^2$.

Figure 5: Vertical hatching. b shows a frontal view of a hatched linear gradient. The dashed lines show the height at which the white and the black layers are laid down. c shows a cross section of the technique used to achieve varying line widths.

3.2 Vertical Hatching

The layer-wise construction of 3D printed objects is used to determine the line pattern used for hatching vertical sections of the 3D meshes. The line widths are not - as one might think - varied by varying the thickness of the layers over different locations; although non-planar FDM printing has seen some developments [Chakraborty et al. 2008], it is not yet available in most slicing software.

Instead, the line width modulation is achieved by using the viscosity of the filament when it departs the nozzle. By extending a given layer past the previous one, the upper layer will sag over the lower one, which results in the lower layer being partially occluded, as has been shown by Reiner et al.. Layers with alternating color use occlusion to modify the visual thickness of lines by applying an amplitude modulated offset to the outlines of each layer. Figure 5 shows the technique and the implementing method.

3.3 Diagonal Hatching

Because objects are built up by discrete planar layers, diagonal surfaces result in layers with differently shaped outlines corresponding to the boundary mesh at horizontal cross sections. In areas of the mesh with angles close to vertical surfaces, consecutive layers tend to be quite similar. When viewed from above, the difference in area between two such consecutive layers would be a thin area that resembles a line. By varying the widths of these area's, hatching can be implemented on diagonal surfaces. The resulting line patterns are therefore concentric polygons of arbitrary shapes.

In diagonal hatching an amplitude modulated offset is applied to the outlines of each layer which is based on the local grayscale value and angle of the surface. The closer the surface angle comes to being horizontal, the further consecutive layers will be apart, which results in a greater distance between the lines of the hatching pattern. In order to make the perceived grayscale value independent of the surface angle, the applied offset is proportional to the distance between the lines. For surface areas with lower slope, the outward offset on the outlines of black layers will be greater at a dark region in the textured mesh. Figure 6 shows the technique used and the method to achieve varying line widths for that hatching technique.

4 VARIABLE OFFSET

This section describes the implementation of the variable offsets required for vertical and diagonal hatching. Before the variable offsets are applied, the model textures are mapped onto the outlines of each layer. Each location on the polygons of the outlines of a layer has a corresponding point in the texture image. The gray scale value at the point in the image is then used to calculate by which distance the outline is displaced at the location.

The displacement for vertical hatching Δ_l is given by equation 1, where C is a constant displacement, L is the texture image lightness value between 0 and 1, h is the thickness of all layers and f is a function such that $f^{-1}(x)$ gives the amount of vertical sagging of a segment overhanging the previous layer by a distance of x . For simplicity, a linear function $f(z) = Az$ is used, where the amplitude A is determined experimentally. Note that the value returned by f is multiplied by $\frac{1}{2}$ because the layer below is offsetted in the opposite direction by a similar amount. The term C is used to shift

Figure 6: Hatching of diagonal surfaces of a sphere with a radial gradient with black on top and white at the equator. The grayscale value decreases linearly with the angle of the altitude. b shows how lines farther toward the middle become more widely spaced. Note how the black lines in the middle are wider than the white lines in the middle and vice versa for the lines near the equator. c shows how varying lines widths are achieved.

the range within which the offset varies with respect to the outlines of the 3D mesh.

The displacement for diagonal hatching Δ_l of a point p in the outlines is given by equation 2, where $\vec{n}$ is the normal vector of the face to which p belongs. Note that the fraction in that equation evaluates to the tangent of the vertical component of the face normal.

$$\Delta_l = C + \frac{1}{2} \operatorname{sgn}(2L - 1) f(h |2L - 1|) \times \begin{cases} +1 & \text{if processing white} \\ -1 & \text{if processing black} \end{cases} \quad (1)$$

$$\Delta_l = C + \frac{1}{2} (2L - 1) h \frac{\vec{n}_z}{\sqrt{\vec{n}_y^2 + \vec{n}_x^2}} \times \begin{cases} +1 & \text{if processing white} \\ -1 & \text{if processing black} \end{cases} \quad (2)$$

The displacement given by the equations above is used to generate a variable offsetted polygon from the existing outlines of a layer. Both vertices and intermediate points on line segments are displaced in the outward direction perpendicular to the polygon. Line segments are offsetted by displacing points sampled at a regular interval, as can be seen in figure 7a.

Given that vertices in the outline generally belong to an edge of the 3D model connecting two faces, which can have different texture image locations at that vertex, applying a variable offset should take two displacement values into account. The offset $\vec{\Delta}_B$ applied to the vertex at the corner between two line segments BA and BC is given by

$$\vec{\Delta}_B = \frac{\Delta_{BA} \vec{BA} |\vec{BC}| + \Delta_{BC} \vec{BC} |\vec{BA}|}{\det \begin{bmatrix} \vec{BA}^T & \vec{BC}^T \end{bmatrix}} \quad (3)$$

, where Δ_{BA} is the offset at B which follows from the texture coordinates at B on the mesh face which AB is located on and likewise for Δ_{BC} .

For inward offsets like the one shown in figure 7b, the offsetted corner bypasses part of the connected line segments. Projecting the

Figure 7: Visual explanation of variable offset applied to parts of a polygon. The original polygon is shown in orange, the resulting offset polygon in black, the sampling points in yellow and the offsets in blue. The grayscale values determine the amplitude of the offset. The green offsets show how variable offsetting of corners should be handled. The red items are omitted from the end result.

offset $\vec{\Delta}_B$ onto the line segments BA and BC gives the distance by which we disregard the sampling points along those line segments. The disregarded sampling points on the edges are not displaced; they are omitted from the variably offset polygon.

When using the above formula, sharp corners which take an outward displacement into account, could result in corners which are displaced by a distance greatly exceeding either displacement value. The resulting corner is therefore beveled when the displacement $\vec{\Delta}_B$ is larger than the bevel distance of both line segments. The bevel distance of a line segment $\vec{BA}$ is given by $b\Delta_{BA}$, where $b > 1.0$ is a constant ratio which determines how much of the corner is beveled off. In figure 7c a bevel ratio $b = 1.1$ was used.

Where line segments in the original outlines are close to each other and the offsets applied are larger than the distance between the line segments, the variably offset polygon would contain self-intersections. The self-intersecting parts of the polygon should be removed by applying a polygon clipping operation which uses a filling rule based on a positive winding number. For further reading, see Vatti [1992].

5 RESULTS

Experiments were performed on an Ultimaker 3 machine, using black and white polylactic acid (PLA) - resp. Ultimaker 9014 and Ultimaker 9016. The nozzle size of this machines is 0.4 mm and a line width of 0.35 mm was used. A default layer height of 0.1 mm was used. Different head movement speeds were used throughout the printing process; most notably, the outer walls were printed with a speed of 15 mm/s. We used a sampling distance of 0.1 mm. Furthermore we applied a static offset of $C = 0.1$ mm for all cases of vertical and/or diagonal hatching.

In an initial testing phase of vertical hatching multiple amplitudes A were tested at 0.01 mm intervals with $f(z) = Az$ in formula 1. After these experiments it seemed that using an amplitude of 0.08 mm yields satisfactory results: 0.08 mm was the minimal tested amplitude at which black pixels in the texture appeared black

and vice versa for white. Results of this tests can be seen in figure 8b. Based on these tests, we estimate that with about 0.16 mm overhang, the filament sags 0.1 mm down.

In tests on diagonal hatching no setting values have been experimented with. While vertical hatching depends on sagging, which in turn depends on quite some physical factors, diagonal hatching has no such dependencies. Figure 8c shows a texture printed at a 45° from a top view.

Tests on horizontal hatching culminated in a result which can be seen in figure 8d, which shows the top of a test print containing a single wall and skin lines. The horizontal hatching method was performed on the top black layer with a fully dense white layer below. The texture was sampled at 0.4 mm intervals. A line distance of 0.7 mm was used, a reference speed of 25 mm/s and a reference line width of 0.35 mm to produce the constant flow of 0.875 mm³/s.

Figure 1 shows hatched objects with geometric and organic shapes and with low and high frequency texture detail, which have been produced using an ad hoc combination of diagonal and vertical hatching. The equations for vertical and diagonal hatching (1 and 2) were combined linearly; the variable offset used adds 95% diagonal hatching to the vertical hatching test setup as described above.

The time it took to print these objects using the 3D hatching techniques was up to 20% longer when compared to printing the same models with a single extruder. This was unforeseen, provided that switching extruder on the Ultimaker 3 typically takes up less than 1% of the print time. We postulate that the difference in printing time is caused by the irregularity introduced by the variable offset. This irregularity causes the print head movement to be limited by its acceleration settings.

6 DISCUSSION

Figure 8b shows a result obtained using hatching on vertical surfaces. The vertical resolution depends on the layer height, which was 0.1 mm. The results obtained using vertical halftoning have a higher horizontal resolution than results obtained by Reiner et al., but also show artifacts most probably due to errors in horizontal positioning and the unpredictability in the nature of sagging. Because vertical hatching is not limited by the horizontal size of dots,

Figure 8: Comparison of a 35 mm × 35 mm surface between the linear halftoning methods presented in this paper. b and c use a layer height of 0.1 mm and d uses a line distance of 0.7 mm. Image courtesy of ©Playboy Magazine 1972.

the resolution of the resulting print is not limited by the line width, which is related to the physical size of the hole in the nozzle.

Moreover, because the surface produced with 3D hatching will be less irregular than when the dithering technique of Reiner et al. is used, we expect the perceived tone to depend less on the azimuth of the viewing angle. When shifting the viewing angle sideways, the perceived grayscale values change less.

Other visual artifacts in vertical and diagonal hatching stem from the fact that a grayscale value based variable offset is applied on a layer-by-layer basis. The perceived grayscale value at a place in a given layer depends on the offsets applied at that place in the layer below and the layer above. Visual artifacts occur where the texture tones or face angles in an area wildly differ between layers. The implementation of vertical and diagonal hatching assume consecutive layers have the same layer thickness. Where this assumption is violated visual defects occur as well.

Horizontal hatching shows a significantly lower resolution than the other techniques presented in figure 8. That’s because the resolution is determined predominantly by the nozzle size, which is high in comparison to the layer height. It should be noted that the hatched print is dark in comparison with the source image. Because the black lines sit on top of a white layer, the perceived tone tends toward black when viewing the surface from a lower altitude angle. Other visual artifacts in horizontal hatching may result from reduced layer bonding when printing thinner lines.

Close examination of figure 8d reveals that at places where the lightness value of the texture is high, black lines appear to be printed thinner, rather than narrower. It seems that these thinner lines only partially block light from the previous layer, resulting in a gray scale value close to the texture lightness even though their width is larger than intended.

7 CONCLUSION

We presented 3D hatching; linear halftoning techniques suited for FDM printing. The results demonstrate the ability to manufacture objects with the appearance of full grayscale textures, while the techniques have little effect on printing time (maximum 20% more than monochrome). Their applicability stands irrespective of the geometry and surface angle and the perceived tone is less sensitive

to viewing angle than existing techniques. Moreover, a higher resolution can be obtained. Another major advantage of this method is that it requires no hardware modifications for dual extrusion FDM.

8 FUTURE WORK

8.1 Optimizing perceived tone with respect to viewing angle

Diagonal hatching has been designed for a viewing angle exactly from above, while one might want to optimize the perceived tone for a specific viewing angle with lower altitude. Alternatively, the perceived tone might be optimized at each place for a viewing angle orthogonal to the local surface. Furthermore, diagonal hatching should in some proportion be combined with either vertical or horizontal hatching, depending on the surface angle. One could apply horizontal hatching to the wall lines themselves, but then diagonal hatching would be affected at higher angles.

8.2 Tone Calibration

The amount of sagging when the filament is extended horizontally over the previous layer determines the amount of occlusion and the perceived tone in vertical hatching. In our tests we used a simple linear formula to compute the required offset to achieve a given amount of sagging. Such a simple formula might mis-render intermediary grayscale values. Once an appropriate model of sagging has been found, the perceived grayscale values of vertical surfaces can be calibrated, and the proportion of vertical and diagonal hatching on diagonal surfaces can be computed.

The formula by which the print speed is computed from the line width in horizontal hatching also provides an opportunity for tone calibration. Given that thin lines seem semi-transparent additional tone calibration might impact the required amount of material per line segment, which affects the line width modulation.

8.3 Colors and multiple materials

One could adopt a similar approach for FDM printers which have more than two extruders; if it is capable of printing with cyan, magenta, yellow, black and white filament, hatching could be used to produce full color prints. This might enable new behaviors on

a macro scale, when applying the offsetting techniques on a meso scale to materials with different properties on a micro scale.

ACKNOWLEDGMENTS

We would like to thank Leo Haslam (Blockade figurine), Beerend Groot (tin can) and COMSOL (connecting rod) for permission to use their models. We would also like to thank Jaime van Kessel for suggestions and improvements.

REFERENCES

- A. Brunton, C. A. Arikian, and P Urban. 2015. Pushing the limits of 3d color printing: Error diffusion with translucent materials. *ACM Transactions on Graphics (TOG)* 35, 1 (2015), 4.
- D Chakraborty, B. A. Reddy, and A. R. Choudhury. 2008. Extruder path generation for Curved Layer Fused Deposition Modeling. *Computer-Aided Design* 40, 2 (2008), 235–243.
- W Cho, E M Sachs, N M Patrikalakis, and D E Troxel. 2003. A dithering algorithm for local composition control with three-dimensional printing. *CAD Computer Aided Design* 35, 9 (2003), 851–867.
- James Corbett. 2012. Reprap colour mixing project. *Final Year MEng Project, Department of Mechanical Engineering, Faculty of Engineering and Design, University of Bath, Bath* (2012).
- Z Corporation. 2005. Z Corporation 3D Printing Technology - Fast, Affordable and Uniquely Versatile. (2005). https://www.uci.ac.cy/arch/documents/3d_Printer_Lab/3D_Printing_Technology.pdf
- Bert Freudenberg, Maic Masuch, and Thomas Strothotte. 2002. Real-time halftoning: a primitive for non-photorealistic shading. In *Rendering Techniques*. 227–232.
- HP. 2014. HP Multi Jet Fusion technology. (2014). h41367.www4.hp.com/campaigns/ga/3dprinting/4AA5-5472ENW.pdf
- Mihaiela Iliescu, Emil Nutu, Kamran Tabeshfar, and Constantin Ispas. 2009. Z Printing Rapid Prototyping Technique and SolidWorks Simulation—Major Tools in New Product Design. In *Proceedings of the 2nd WSEAS International Conference on Sensors, and Signals and Visualization, Imaging and Simulation and Materials Science*. World Scientific and Engineering Academy and Society (WSEAS), 148–153.
- Tim Kuipers. 2017. CuraEngine: 3D Hatching: Linear Halftoning. (2017). DOI: <http://dx.doi.org/10.5281/zenodo.571229>
- Irving Lavin. 2004. Claude Mellan's Holy Face. (2004).
- Qun Lou and Peter Stucki. 1998. Fundamentals of 3D Halftoning. In *Electronic Publishing, Artistic Imaging, and Digital Typography*. Springer Berlin Heidelberg, 224–239.
- Mcor. 2013. How Paper-based 3D Printing Works: The Technology and Advantages. (2013). <http://mcortechtechnologies.com/wp-content/uploads/2013/04/MCOR-WP-19032013-EU> (accessed April, 2017).
- Claude Mellan. 1649. The Sudarium. (1649). <http://hdl.handle.net/10934/RM0001.collect.152641>
- Emil Praun, Hugues Hoppe, Matthew Webb, and Adam Finkelstein. 2001. Real-time hatching. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*. ACM, 581.
- Tim Reiner, Nathan Carr, Radomir Měch, Ondřej Št'ava, Carsten Dachsbacher, and Gavin Miller. 2014. Dual-color mixing for fused deposition modeling printers. In *Computer Graphics Forum*, Vol. 33. 479–486.
- Stratasys. 2016. Stratasys J750. (2016). http://usglobalimages.stratasys.com/Main/Files/Machine_Spec_Sheets/PSS_PJ_StratasysJ750_0217a_Web.pdf
- Bala R Vatti. 1992. A generic solution to polygon clipping. *Commun. ACM* 35, 7 (1992), 56–63.
- Kiril Vidim, Szu-Po Wang, and Jonathan Ragan-kelley. 2013. OpenFab: A Programmable Pipeline for Multi-Material Fabrication. *ACM Transactions on Graphics* 32, 4 (2013), 1–11.
- Akio Yamamoto. 2006. Producing engraving-style halftone images. (2006). US Patent 7,126,723.