

Technical Perspective

From Virtual Worlds to Digital Fabrication

By Sylvain Lefebvre

COMPUTER GRAPHICS HAS enabled game developers to create vast open worlds full of wonders and dangers that players routinely explore and interact with. These worlds are not made of physical matter; instead, they are defined by a surface geometry—often triangles—and material parameters, such as color and reflectance. Rendering algorithms use this information to form a rapid sequence of images on screen, simulating the viewpoint of a moving virtual observer.

This ability to depict large open worlds has been a fundamental challenge from the infancy of computer graphics. The problem is twofold: describing the virtual world in all its intricate details, and being able to store and process this data for rendering. Describing a small virtual scene like a single room is not a problem: a skilled artist can use CAD software to model the surface geometry of each object, as well as specify the color of its parts. However, this quickly becomes impractical as the virtual scene grows in extent to become, for instance, an entire planet. Similarly, even if an entire planet could be modeled, the data would simply not fit the computer memory.

For these reasons, the field of computer graphics has focused intensely on dealing with very large amounts of geometry and material information. This has had a deep influence on the software, hardware, and industrial practices in our field. One key idea, pioneered by researchers such as Ken Perlin and Kenton Musgrave, is the notion of *proceduralism*: The idea that detailed geometric and material information does not have to entirely exist in memory. Instead, it can be generated on-the-fly, when needed, for rendering a single viewpoint. This is a powerful idea: The screen onto which an image is displayed has a limited resolution, and the amount of data visible in a single viewpoint is only a


tiny subset of what could be the entire universe. By computing, or rather *synthesizing* only the content required for the current view, it becomes possible to explore worlds without bounds: more content is synthesized as the user wanders deeper into the virtual landscape. The content is then described by procedures, small algorithms that generate details whenever required from a simpler description of the scene.

How does this relate to additive manufacturing? As noted in the following paper by Vidimče et al., the rapid increase in both print resolution and print volumes, combined with the ability to mix different materials, leads to a very similar situation. Describing a 3D model in all its intricate details becomes rapidly infeasible. The challenge is not only in describing the object using available tools, but also in being able to store and process this description before fabrication. It might seem surprising: this is, after all, a single object. However, 3D printing requires specifying the material *at every point in the volume*, while most often virtual worlds require only describing surfaces. In addition, the print resolution is increasing toward micron accuracy. As

3D printing requires specifying the material at every point in the volume, while most often virtual worlds require only describing surfaces.

a consequence, the number of points that must be specified to fully exploit the capabilities of the 3D printers grows very rapidly.

Additive manufacturing technologies fabricate an object layer after layer, from bottom to top. Each layer can be thought of as a two-dimensional grid of little cubes, where each cube is either empty, or will be filled with one or a mixture of materials. Taken together, the layers form a large three-dimensional grid of cubes, called *voxels*. Even with today's limitations, a print using the full extent of the printer can have billions of voxels. Fortunately, when the part is being fabricated, only a single layer needs to be in memory; this is akin to the limited viewpoint of the virtual observer in a virtual world. Thus, the same methodologies apply—rather than storing the object in all its intricate details, the details can be synthesized only when needed by the fabrication process, layer by layer.

The authors of OpenFab propose to revisit the processing pipeline that turns a 3D model into machine instructions in light of the solutions developed in computer graphics. In particular, rather than specifying the object by handpicking a material in each of its voxels, users can write small algorithms that synthesize the content when it is needed for fabrication. This integrates in an elegant pipeline that affords for unprecedented design flexibility, while simultaneously answering computational challenges. Suddenly, it becomes possible to fully exploit the high-resolution capabilities of the additive manufacturing processes. This unlocks a vast number of possibilities, from aesthetics to novel optical and mechanical properties, triggered by micro-structures embedded within the object's volume. 

Sylvain Lefebvre (sylvain.lefebvre@inria.fr) is a senior researcher at Inria, Nancy-Grand Est, France.

Copyright held by author.