SYSTEMS ISSUES IN SOLID FREEFORM FABRICATION

Michael J. Wozny
Rensselaer Design Research Center
Rensselaer Polytechnic Institute
Troy, New York 12180

Abstract

This paper is concerned with the systems aspects of the Solid Freeform Fabrication (SFF) technology, i.e., the issues that deal with getting an external geometric CAD model to automatically control the physical layering fabrication process as directly as possible, regardless of the source of the model. The general systems issues are described, the state of systems research is given, and open research questions are posed.

1. INTRODUCTION

Prototyping requirements today, and general manufacturing requirements in the future, call for the rapid, fabrication of one-of-a-kind, structural strength parts driven directly by computer-sensible geometric data. The general class of Solid Freeform Fabrication (SFF) systems, where material is added layer by layer, has the potential to become the dominant prototyping technology in the near term, and even a key manufacturing technology in the future.

Solid Freeform Fabrication deals with the problem of fabricating, under computer control, a CAD description of a desired part by selectively solidifying or bonding one or more raw materials into a thin layer, representing a horizontal slice of the desired part; and then fusing the successive thin layers into a 3D solid object. The material may be a gas, liquid, powder or a thin solid sheet, while the solidification process may be polymerization, sintering, chemical reaction, plasma spraying or gluing. The geometry data may be CAD geometry, CAT or MRI imaging data, or special forms like contour slices. Only CAD geometry is assumed in this paper, since it is the dominant source of data for SFF today.

General references on various aspects of the technology are [Marcus, et. al. 91, Bjorke 91, Kruth 91, Arline 91, Bourell et. al. 90, Deckard 86]. Other references can be found in the proceedings of the annual symposia on Solid Freeform Fabrication at the University of Texas at Austin, and the annual conferences at the University of Dayton (Ohio). Current events are published in the monthly newsletter [Cohen].

A high-level systems view of the overall SFF process is shown in Fig. 1.1. The computer-sensible geometric CAD model describes the desired final shape of the physical part. The information processing subsystem converts this input geometric model into a form suitable for controlling the solidification process. The physical solidification process subsystem creates the actual physical part. (The term solidification is used because it is the dominant processing technology today.) The physical process parameters depend upon the characteristics of the material(s), the solidification process, the environment in the solidification chamber, as well as precision, fabrication speed, and geometric shape.

As we shall see, the control strategy is basically feedforward control, since modeling of the solidification process is still a very active research area, and adequate sensor technology is lacking. Most progress in research has been in the manipulation of geometric shapes.

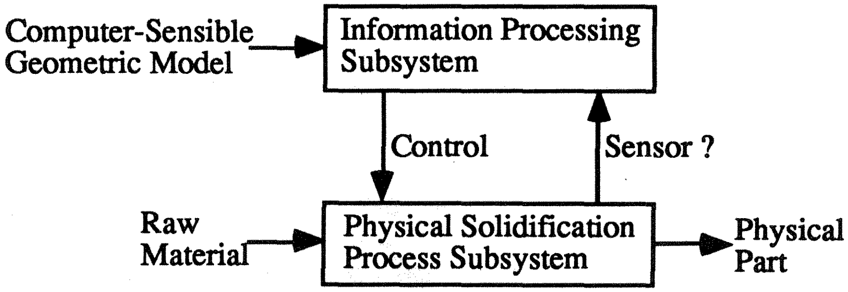

Figure 1.1 Major Subsystems

2. PHYSICAL LIMITATIONS ON SHAPE

The input geometric model represents the desired shape of the physical part. In order to develop an appropriate control strategy the ideal geometry must first be sliced into layers.

2.1 Slices of Finite Thickness. Since the layers of the physical material must have finite thickness, we are faced with the volume sampling problem shown in Fig. 2.1. This 3D aliasing effect limits the accuracy of the final physical part.

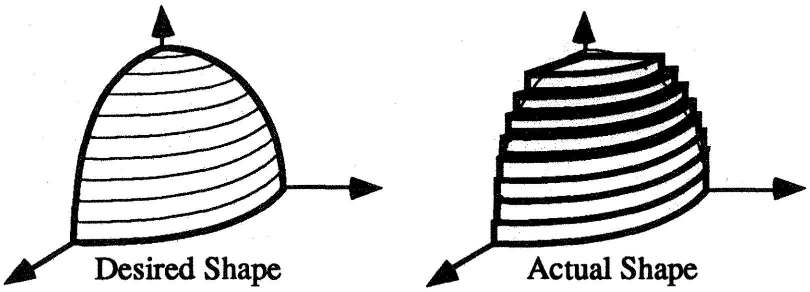

Fig. 2.1 Finite Layers

From a systems point of view, layer thickness is primarily influenced by the trade-off shown in Fig. 2.2. However, physical characteristics such as depth of solidification, density and desired strength are also key considerations.

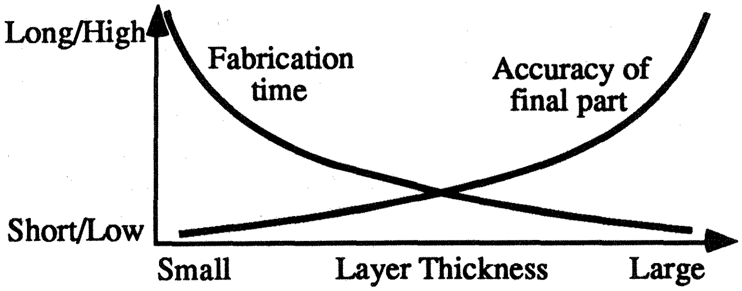

Fig. 2.2 Trade-off in Layer Thickness

One means for reducing the effect of volume sampling, for layers of constant thickness, is to choose the extent of each layer so that it best fits the enclosed solid slice. Fig. 2.3 illustrates one approach for a solid of revolution where the radius of each layer depends on the ratio of the layer volume to the part volume. Criteria other than equal volumes can be chosen. Also, the calculation becomes much more difficult for parts which are not solids of rotation. This subject requires further study.

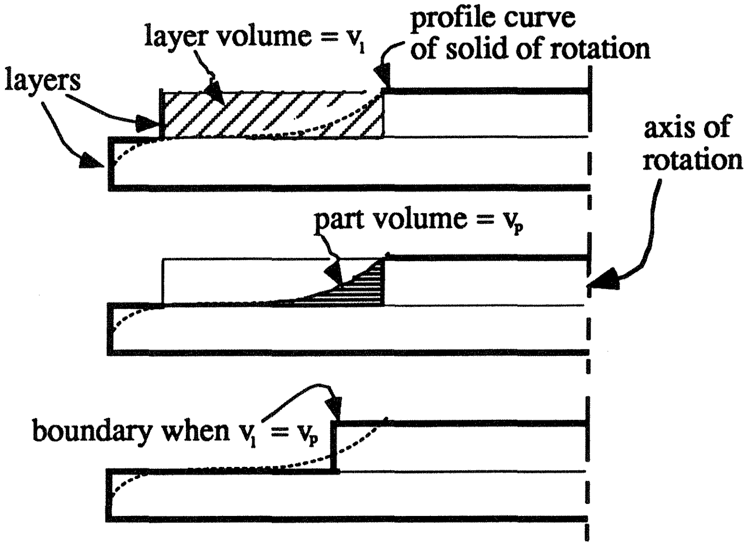

Fig. 2.3 Reducing Volume Aliasing Effects in Solids of Rotation

In a recent study, 3D anti-aliasing can also be achieved by varying thickness of the slice (sometimes called adaptive slicing). The thickness of the slice has been related to local surface curvature, number of contours in each slice, and distance between line segments in different slices [Dolenc et. al. 92].

The above method contributes to the reduction of the 3D aliasing, but does not eliminate it. Aliasing can be further reduced only by taking advantage of physical phenomena.

2.2 Scan Vectors of Finite Thickness. There is also an aliasing effect due to the finite thickness of scan lines on each layer, as shown in Fig. 2.4. From a systems point of view, compensation techniques such as overlapping and multi-directional scan lines, and tracing the final boundary can be used to minimize the effects of sampling. (Scanning strategies for efficiency are discussed in Section 9.) However, the ultimate accuracy depends on the beam size and beam energy cross-section. These and other sources of error were studied by [Bjorke 91].

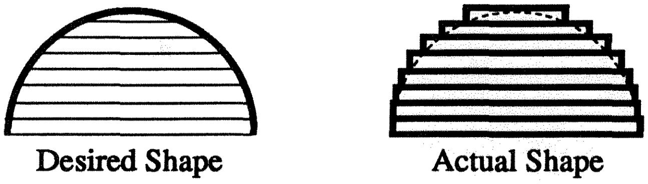

Fig. 2.4 Scan Line Aliasing

The next section describes cases where apparently harmless solid models and modeling techniques, severely impact the fabrication of physical parts.

3. PHYSICAL IMPLICATIONS OF CAD SOLID MODELS

Regardless of field, models are developed to answer specific questions. If the models are extended to new contexts, then all assumptions must be carefully re-examined. This is the case with CAD geometry models which now are being extended to directly drive the layered fabrication process. Solid models created for purposes of visualization or analysis are not necessarily appropriate for SFF.

3.1 Coincident Surfaces. The following example illustrates a problem which occurs when a solid model with coincident surfaces is used to fabricate a part [Aubin 92]. Such a model is perfectly adequate for visualization or volume analysis, but fails for SFF.

Suppose a solid model is created in which the coincident (inner) surfaces of adjoining constituent solid elements are not removed. This tactic is commonly used when transitioning between different solid sections. The resulting fabricated physical object may exhibit deformations or poor strength characteristics in the region of the coincident surface. Fig. 3.1 illustrates this problem.

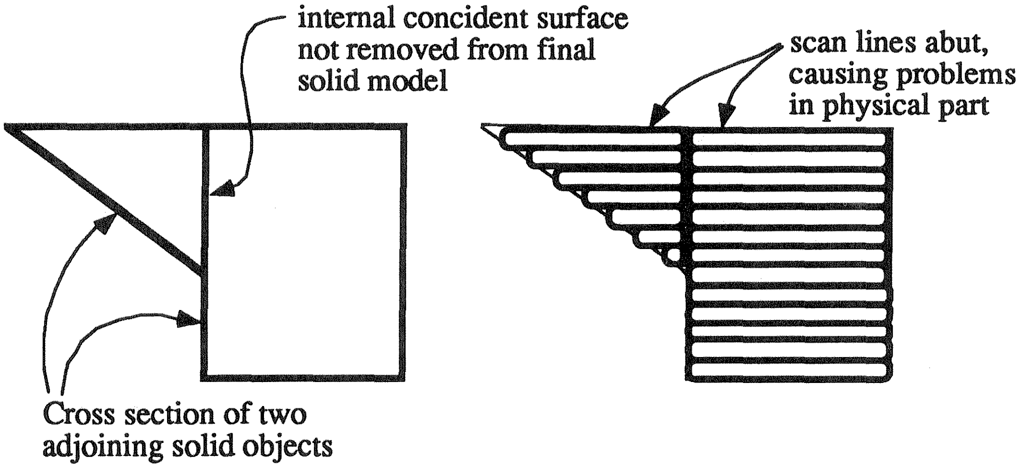

Fig. 3.1 Concident Surfaces

What other situations exist where improperly (for SFF) generated CAD models may result in undesired physical characteristics? In another case, abutting solids can also result in non-manifold topology [Weiler 86] conditions where there is ambiguity of what constitutes the interior or exterior of the final object [Rock 91, Bøhn 92 & Wozny 92]. Such cases need to be identified and studied. We must understand the resulting physical manifestation implicit in the geometric model. A rule base may be needed to evaluate troublesome cases when validating solid models (more later).

Next, suppose the user discovers an anomaly in the geometric solid while fabricating a physical part. It would be useful to send this anomalous data back to the CAD system and modify the existing geometric data without having to return to the original model in the CAD system. What associativity information is needed so that models which have been already sliced and scanned can easily be modified in a global sense? This is also an open question.

3.2 Postprocess geometry in the CAD System. The situation in the previous paragraph leads to the question of where should the processing of solid models into scan lines take place? One could postprocess in the CAD system, geometric part data directly into scan vectors suitable to control the physical fabrication process. The CAD system has a common geometric representation and, therefore, can generate exact slices and scan vectors.

Although this approach may be efficient for a specific CAD/SFF device combination, transferring scan vectors has its drawbacks. The approach requires extensive low level knowledge about the SFF device, such as internal timing and control mechanisms. Errors can be very costly, since there is minimal error protection and recovery. The scan line files will be extremely large for real problems. Verification of the data would not be easy, as well as compensating for SFF idiosyncrasies such as shrinkage.

This situation is similar to plotting a curve on an XY-plotter, where the curve is subdivided into short straight line segments in the CAD system, and then the segments are sent to the plotter in the proper sequence and timing. Using another analogy, it's like programming in assembly language rather than a high level language like C or Pascal. Although it may work very well in specialized situations, it does not appear to be strategic. The high level descriptions change most slowly and remain more modularized than low level data and descriptions.

3.3 Processing Geometry in the SFF Device. The strategic approach is to accept high level 3D geometry descriptions from the CAD environment and perform all subsequent slicing, scanning and other related processing in the SFF device. This approach, in principle, allows a general fabrication capability which can accept data from any CAD system, making the SFF device independent of the CAD environment.

High level geometric descriptions permit general procedures for the validation, orientation and nesting of parts, increasing accuracy and machine utilization. Questions such as part placement for optimal fabrication including part build time, scanning efficiency, part surface finish, and accuracy on critical dimensions need to be investigated. Such investigations will likely find ties to other research in tolerancing and feature-based modeling. Since most parts use only a portion of the maximum allowable part fabrication volume, they can be nested to increase machine thruput. This nesting problem requires nontrivial extensions of two-dimensional nesting concepts to three-dimensions. Physical process properties may also be incorporated in such algorithms to account for part shrinkage and proximity to adjacent parts.

If such operations are to be performed in the SFF device, then the device must maintain maximum flexibility by accepting a high level geometric description from any CAD system, regardless of math form and representation. Attacking the problem of proliferating of math forms is discussed next.

4. COMMON MATHEMATICAL FORM

Unfortunately, many different mathematical forms have been developed to model the extensive range of part geometries A subset of these forms is illustrated in Fig. 4.1. Each mathematical form requires its own set of slicing, scanning, and other algorithms, and there-in lies the problem. It is impractical for SFF devices to support this range of input models.

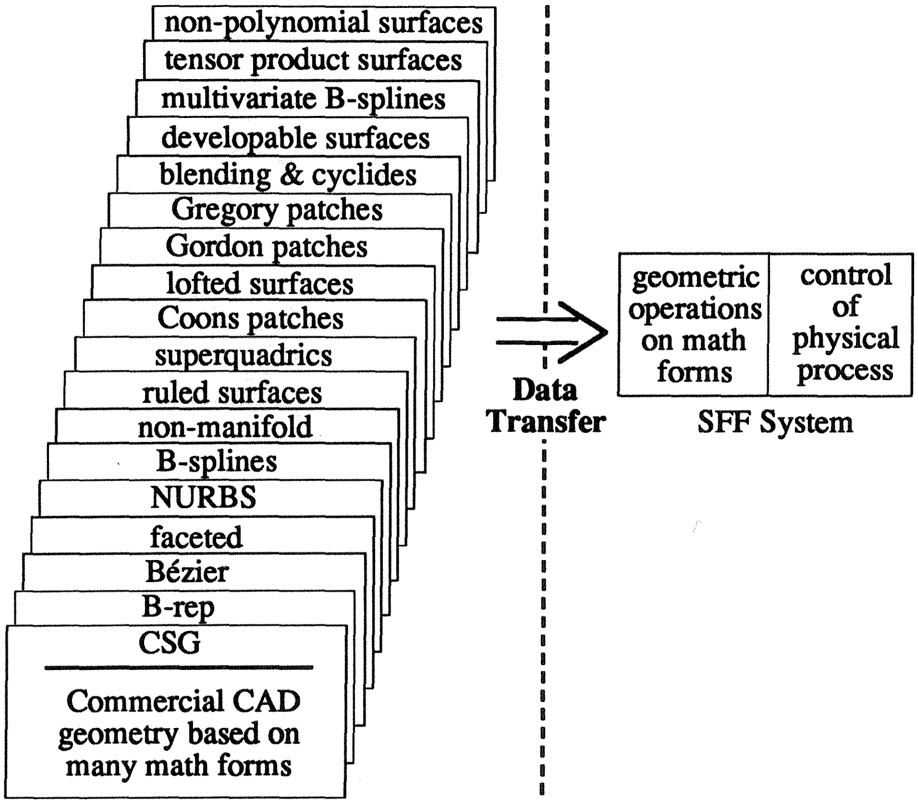

Fig. 4.1 Many Math Forms for Geometric CAD Parts

The strategic approach is to settle on one, or at most a few, common mathematical forms and require that all geometric models be approximated with the chosen common forms. Accepting this premise, the next question is, which forms are the most appropriate? Should the simplest possible form of planar triangular facets be chosen, or a reasonably general form such as non-uniform rational B-splines (NURBS)? Clearly, using a low degree polynomial surface element to approximate a given geometric model will require more elements to achieve a given precision, than would a high degree polynomial surface element. On the other hand, low degree elements involve simpler algorithms for slicing and scanning, implying faster computation. A study of these trade-offs, illustrated in Fig. 4.2, for a class of representative part models would be enlightening.

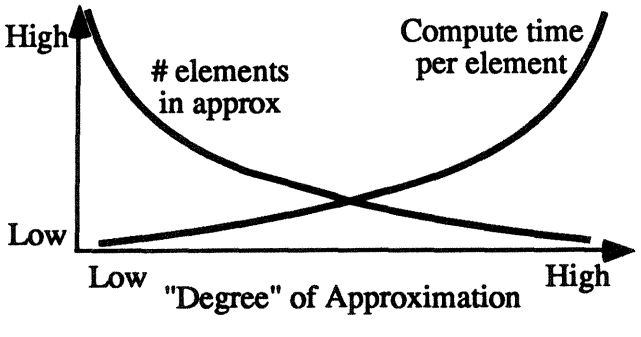

Fig. 4.2 Trade-off in Compute Time and Complexity

Although a NURBS surface element (patch) is relatively general, and can represent the majority of real parts, it could be overkill if most of those parts consist primarily of simple geometries, implying unnecessary computations, and overhead.

Today the de facto industry standard [Stereolithography Interface Specification] is faceted representations, i.e., the approximating surface element is a planar triangle. Unfortunately, the number of triangular facets needed to maintain high precision in complex parts becomes prohibitive. Higher degree polynomial and rational polynomial approximations are necessary to maintain precision. Fortunately, a precedent has been set in the evolving CAD data exchange standards, IGES, and STEP ["ISO CD 10303 - 42]. Since the purpose of these evolving standards is to exchange CAD data (e.g. geometry) among different CAD environments, as well as applications, without revealing proprietary internal data representations, then SFF devices should also accept these geometric forms, which includes NURBS surfaces. Consequently, slicing and scanning algorithms need to be developed only for the STEP math forms. Although the general mathematics for general surface-surface intersections, including NURBS, has been developed, no reliable experimental package is widely available.

A note of caution. If a CAD model was originally created entirely with NURBS surfaces, then the above CAD exchange standards guarantee that the model can be sent to a SFF device essentially intact. However, if a geometric CAD model is of a different math form which must be approximated with NURBS patches, including continuity conditions, then the problem becomes extremely difficult. Subsequent sections will show that even the simple planar patch approximation spawns a whole host of nontrivial problems. All of these problems become orders of magnitude more difficult for NURBS patch approximations. The advantage of NURBS is that most models will be originally generated in this representation, eliminating the need for approximations.

The next section examines some of the problems which arise when creating facet approximations.

5. FACET APPROXIMATION

Unfortunately, many commercial tessellation algorithms used by CAD vendors today are not robust, creating polygonal approximation models having: gaps (cracks, holes, punctures), i.e., missing facets; incorrect or inconsistent normals; non-manifold topology [Weiler 86] conditions, where three or more facets share a common edge or two solids are tangent along a common boundary; edge and point degeneracies consisting of equal or collinear edges; and self-intersections [Bøhn & Wozny 92]. The underlying problem is due, in part, to the difficulties in tessellating trimmed surfaces, surface intersections, and controlling numerical error. A surface intersection anomaly which results in a gap is shown in Fig. 5.1.

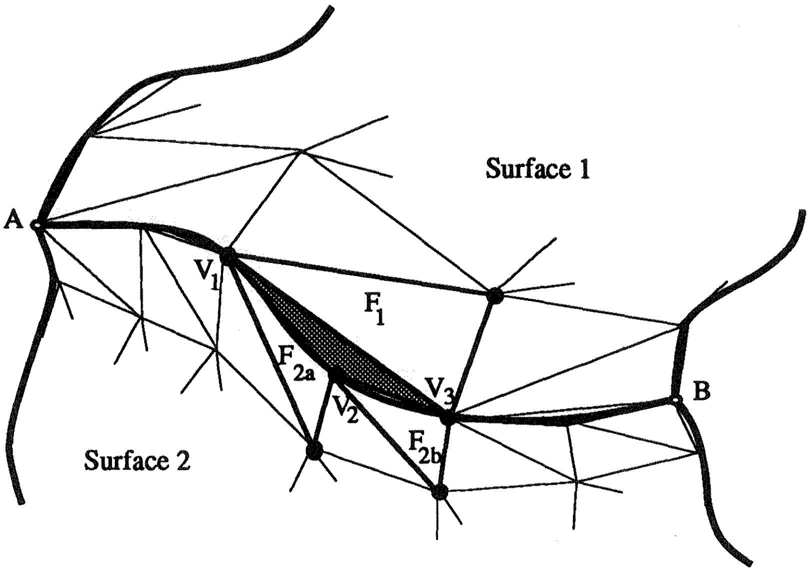

Fig. 5.1 Gaps Due to Missing Facets

The missing facet in the geometry model means that the SFF device has no defined stopping boundary on a given slice, and processes material right to the physical limit of the device, creating a stray physical solid line and ruining the part. This is an example of an invalid model.

Although a number of anomalies in tessellated models have been identified, they are by no means well understood. The more interesting problem, described next, is how to repair these invalid models.

6. VALIDATION OF GEOMETRY AND REPAIR

A basic requirement of the CAD geometric model is that it must realize a valid physical part from the SFF process. A model which meets this requirement is called valid. How does one ensure, *á priori*, that the CAD model is valid? The first step in validation deals with having a closed shell, i.e., no missing facets. (Model validation can be described in terms more general than tessellated models, for example, NURBS.)

If the model is invalid, then procedures must be developed to repair it. If a tessellated model is found to have gaps, then it must be repaired, i.e., the gaps must be filled with a "suitable" approximation of triangular facets. Since we only have a sampled surface available, i.e., the model is "correct" only at discrete points All the data which defined the original surface is not available. In addition, if some of the discrete points are also missing, then even the original faceted surface cannot be restored.

The model validation and repair problem for tessellated models can be stated as follows: Given a facet model, i.e., a set of triangles defined by their vertices, in which there are gaps, i.e., missing one or more sets of polygons, generate a "suitable" triangular surface which "fills" the gaps.

Preliminary research has shown the repair problem is difficult and not at all obvious [Rock 91]. Non-manifold edges must be resolved such that each facet has only one neighboring facet along each edge, i.e., reconstructing a topologically manifold surface. The problem of cracks requires the identification of the bounding edges and their sets of closed hole-boundary loops. Vertices can be eliminated from this loop by forming new triangular facets with neighboring vertices until the hole-boundary disappears. Finally, the problem of possible self-intersections resulting from the earlier (possibly numerically imprecise) operations must be resolved.

An integral part of the above solution is to ensure a correct surface orientation for all facets. If the original facet orientations cannot be trusted, then it cannot be assured that a CAD-model with internal voids and solids can be correctly repaired. See [Bøhn & Wozny 92] for recent results on the repair problem.

7. EFFICIENT SLICING

After validating and conditioning, the geometric model is first sliced into layers, and each layer converted into scan lines as shown in Fig. 7.3. The scan lines determine the toggle points (on/off points) for the laser beam controller. The slicing or surface-plane intersection algorithms are very computation intensive. As a result it is important to make these algorithms as efficient as possible.

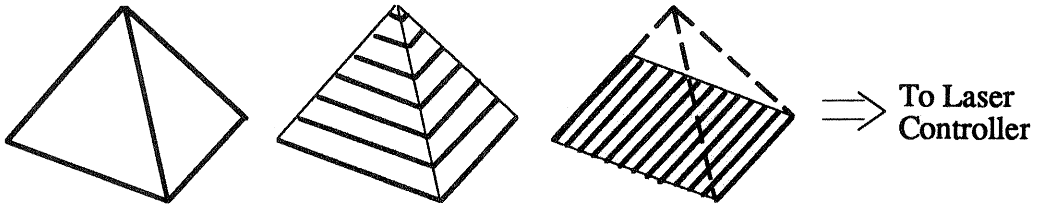

3Dmodel

Slice 3D model
into
2D planar layers

Convert each 2D layer
into
1D scan lines

Fig. 7.1 Geometric Slicing Operations to Create Scan Lines.

To gain efficiency, topology information about the facets is needed. Using topology to incrementally intersect each triangle at a given vertical level as the algorithm marched around the facets proved to be a very efficient technique. Details are given in [Rock 91, Rock & Wozny 91b].

Other techniques, such as ray-tracing, were aimed at generating scan lines directly from the 3D model, but in general, such methods proved to be very computation intensive and thus slower [Rock 91]. Slicing NURBS surfaces is significantly more computation intensive than the polygonal marching algorithm above, as expected.

8. FILE-FORMATS

The current *de facto* SFF input file-format standard, STL, was developed by 3D Systems, Inc. [Stereolithography Interface Specification]. It consists of an unordered list of triangular facets without any topological information other than the orientation of each facet (i.e., which side of the facet is the material-side). Consequently, it is not always obvious how to mate neighboring facets, nor is it a trivial matter to determine which facets are neighbors due to numerical imprecision.

A majority of the problems encountered with the STL file format is the lack of topological information. The redundancy in storing duplicate vertices and edges is shown in Fig. 8.1. An algorithm for generating topological information is given in [Rock 91, Rock & Wozny 92].

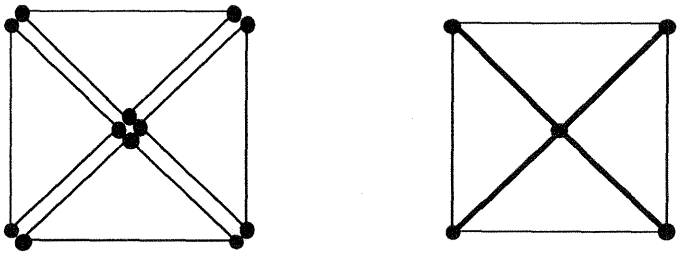

STL Format: 12 Verticies
(36 doubles)

RPI Format: 5 Verticies
(15 doubles, 12 pointers)

Fig. 8.1 Edge and Vertex Redundancy in STL Format

A new file format, called the RPI format, [Rock 91, Rock and Wozny 91a] is significantly more compact than the STL format, eliminates the redundancy in STL, maintains topological information, and simplifies the task of ensuring that a model is valid, i.e., no missing facets, etc. The RPI format is derivable from STL format data. It is extensible, represents facet solids, and includes information about facet topology. Topological information is maintained by representing each facet solid entity with indexed lists of vertices, edges, and faces. Instead of explicitly specifying the vertex coordinates for each facet, a facet can reference them by index number, reducing redundancy. The RPI format file is composed of a collection of entities, each of which internally defines the data it contains, and conforms to the syntax defined in the syntax diagram shown in Fig. 8.2.

Each entity is composed of an entity name, record count, schema definition, schema termination symbol, and the corresponding data. The data is logically subdivided into records which are made up of fields. Each record corresponds to the definition provided by the schema. Each field corresponds to one variable type in the Type Definition. Entity definitions have been developed for specifying facet solids, CSG solids, operations and transformations on these solids, as well as process specific data [Rock 91].

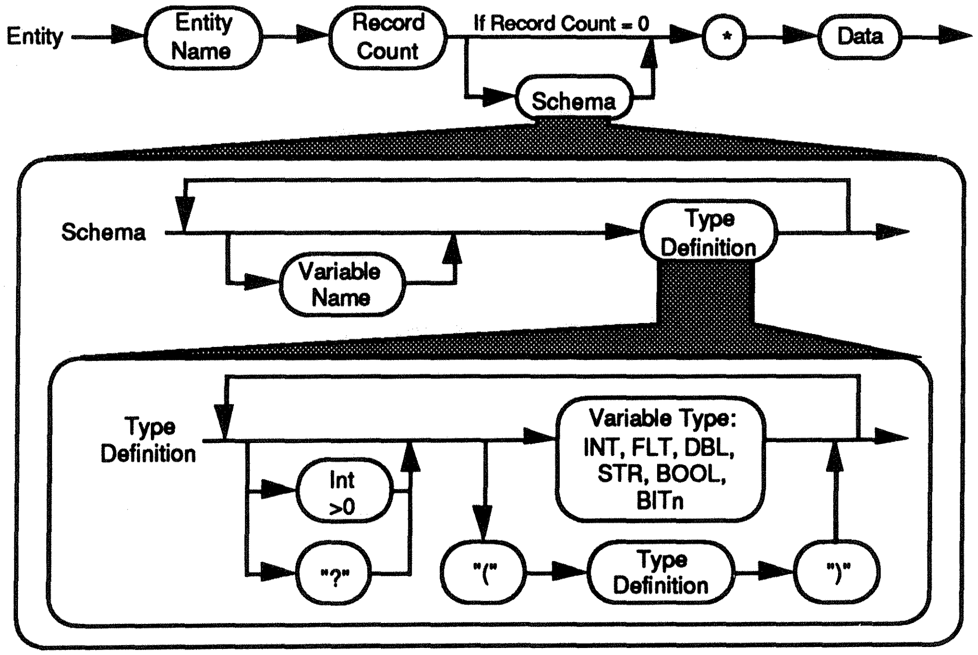

Fig. 8.2 RPI Format Entity Syntax Diagram, from [Rock 91]

9. INTELLIGENT SCANNING

It is not clear that the standard raster scan is the best scanning strategy for SFF. For example, experiments have shown that tracing the boundary of a part provides a better edge definition. This section, based on preliminary research by [Sankauratri 92], describes several strategies that are derived from the geometric shape of a polygonal slice.

9.1 Longest Edge Scan. In this case the scanning is done parallel to the longest edge of the polygon. Fig. 9.1 illustrates the advantage of this strategy. It results in longer uninterrupted active periods of processing material.

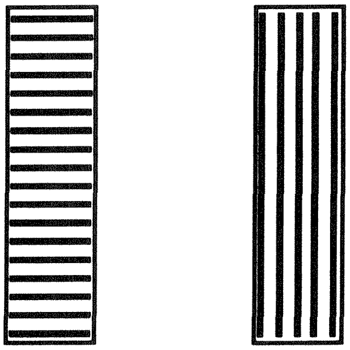

Raster Scan

Longest Edge Scan

Fig. 9.1 Advantage of Longest Edge Scanning

This strategy can be achieved simply by rotating the physical part or the laser to achieve the proper orientation, and raster scanning.

9.2 Adaptive Longest Edge Scan. In this case the scan proceeds perpendicular to the longest edge until a transition (i.e., a vertex) is reached. Then all the remaining edges, including the new ones formed during the scanning are searched to find the current longest edge. The scan proceeds perpendicular to the new longest edge until another transition is reached. This approach, illustrated in Fig. 9.2, minimizes the number of toggle points.

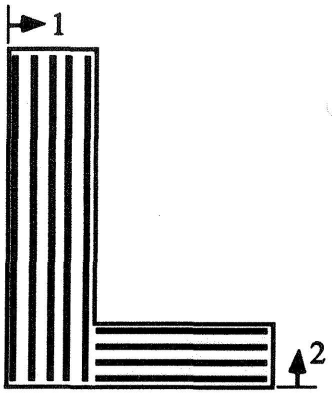

Fig. 9.2 Adaptive Longest Edge Scanning

In this approach and the following one, the changes in scanning direction will effect the physical characteristics of the part. This aspect requires further investigation.

9.3 Spiral Scan. In spiral scanning, the center of the spiral is placed at the centroid of the polygonal slice. Equi-angular rays are extended from this point. The scanning is incrementally generated from one ray to the next. An advantage of this approach is that the part is fabricated from the inside toward the outer boundary. This allows expansion due to the heat of the fabrication process to constantly move toward the unprocessed material. It appears that more accurate parts could be obtained in this manner. From a geometrical point of view, spiral scanning produces one long scan line for convex polygons.

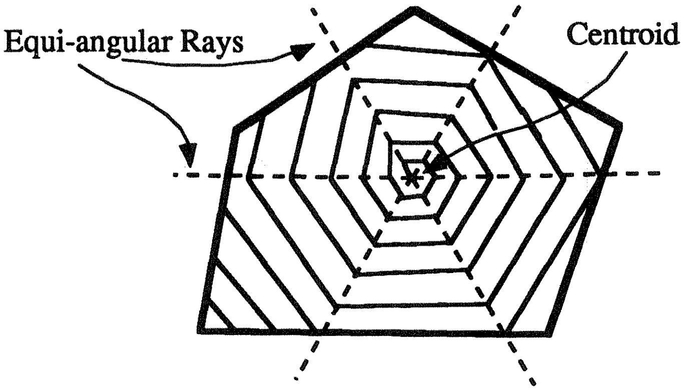

Fig. 9.3 Spiral Scanning

9.4 Medial Axis as a Scanning Strategy. If one extends spiral scanning to polygons with holes or to scanning about internal lines rather than centroids, the analogy to the medial axis [Prinz] is uncanny. Geometrically, the medial axis is the locus of centers of circles of various radii whose circumferences touch the object boundary at two points. Intuitively, of one thinks of the 2D polygon as an area which burns (yes, fire) uniformly, then the medial axis is the set of lines to which the object eventually burns when the boundary (including hole boundaries) is set on fire. The medial axis for a rectangular slice is shown in Fig. 9.4.

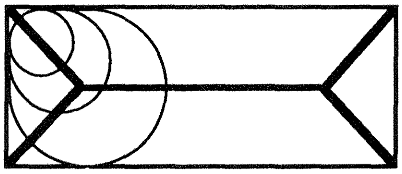

Fig. 9.4 Medial Axis for Rectangle

Unfortunately, the medial axis itself is not symmetrical enough to generate a continuously increasing spiral with equi-angular rays. But the concept does point toward a type of symmetry that is needed to support this type of scanning strategy. More research is needed to develop a uniform theory of scanning strategies, based on object symmetries.

10. INTELLIGENT CONTROL

For reasons of specificity, this section is concerned with the laser sintering process. The last section described an intelligent scanning strategy for controlling, for example, the laser beam. However, the overall goal is not simply intelligent scanning, but intelligent sintering. Intelligent sintering closes the feedback loop around the entire physical process, as illustrated in Fig. 10.1.

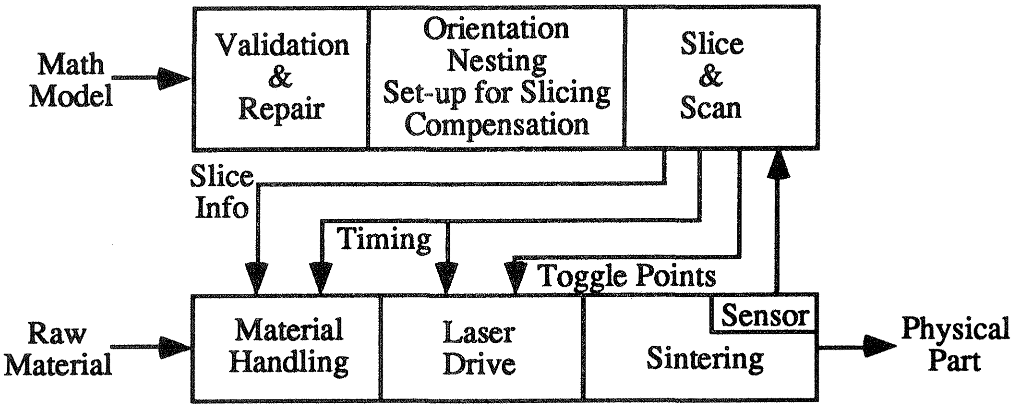

Fig.10.1 Block Diagram Showing Feedback

Research issues consist of modeling the physical process as the plant dynamics in the control loop, and developing the appropriate sensor technology to feed back the proper signals.

11. CONCLUSIONS AND FUTURE TRENDS

Although the technology is still in its infancy, it has the potential to be a dominant technology in data driven rapid prototyping, and ultimately in data driven fast, flexible, lot-size-of-one, manufacturing.

The major systems problem today deals with data transfer, namely, what should be the standard interface specification for the SFF system so that it can accept geometric data from any CAD system. The adoption of the STL format represented an appropriate approach in the early stages of the commercial technology in 1988, but now needs a major revision. The RPI format is a big step in this new direction. It is also clear that the technology must move beyond tessellated geometric models and deal with precise geometry. The evolving ISO STEP/PDES data exchange standard, especially the geometrical aspects, will have a major role in determining the data interface of future SFF systems.

The next major thrust will most likely be in efficient slicing and scan conversion for precise models. Questions which deal with part orientation, nesting, and intelligent scanning, have yet to be addressed in the open literature.

SFF systems, when viewed as lot-size-of-one manufacturing machines will provide an advanced capability where manufacturing, design, and materials all come together into a meaningful whole. This will change our design methodologies as well as eliminate or radically change traditional manufacturing procedures, such as, planning, tooling, and fixturing.

This work focused on systems and geometric issues. But it is clear that in the future CAD will extend beyond the predominate geometric volume considerations, and integrate more material characteristics. Instead of building strength into parts by adding bulk (geometric volume), SFF allows the possibility of changing material properties to achieve strength. One can also consider changing surface material to increase wear characteristics. The issue of blending from one material to another needs to be investigated, as well as a range of new applications, such as smart materials and devices developed in layers. Finally, we need to incorporate material characteristics into our CAD systems and download such characteristics directly to SFF machines.

12. ACKNOWLEDGMENTS

The author thanks his students Stephen Rock and Jan Helge Bøhn, both pursuing Ph.D. research in this area, and Sridhar Sankauratri, who is studying intelligent scanning for his MS degree, for many fruitful discussions and collaboration. Many of the results described here can be found in their theses. He also thanks Dick Aubin, Pratt & Whitney (a division of United Technologies) for providing a number of industrial STL files, and a lot of insight and experience.

This research was supported by the National Science Foundation Strategic Manufacturing Initiative Grants No. DDM-8914172 (Rensselaer Polytechnic Institute) and DDM-8914212 (University of Texas at Austin), the Office of Naval Research, the New York State Center for Advanced Technology, and the RDRC Industrial Associates Program.

Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author and do not necessarily reflect the views of the National Science Foundation, or any of the industrial sponsors.

13. REFERENCES

Arline, J., "A User Survey of Problems with Rapid Prototyping Systems," *Rapid Prototyping Report,* vol. 1, no. 6, November 1991, pp. 7-8.

Aubin, Dick, Private Communication, May, 1992.

Bjorke, O., "How to Make Stereolithography into a Practical Tool for Tool Production," *Annals of the CIRP*, vol. 40/1/, 1991.

Bøhn, J. H. and M.J. Wozny, "Automatic CAD-Model Repair: Shell-Closure," *Solid Freeform Fabrication Symposium 1992*, Published by the University of Texas at Austin, 1992.

Bourell, D.L., Beaman, J.J., Marcus, H.L., and J.W. Barlow, "Solid Freeform Fabrication An Advanced Manufacturing Approach," *Solid Freeform Fabrication Symposium 1990*, Published by the University of Texas at Austin, 1990.

Cohen, A.L., eds., *Rapid Prototyping Report.*, The newsletter of the desktop manufacturing industry, CAD/CAM Publishing, Inc., 841 Turquoise St., Suites D&E, San Diego, CA 92109-1159.

Deckard, C.R., "Part Generation by Layerwise Selective Sintering," M.S. Thesis, Department of Mechanical Engineering, University of Texas at Austin, 1986.

Dolenc, A., Mäkelä, I., and R. Hovtun, "Better Software for Rapid Prototyping with INSTANTCAM", IFIP Transactions B. *Human Aspects in Computer Integrated Manufacturing*, G.J. Olling and F. Kimura, eds. Elsevier Science Publishers B.V. (North-Holland), 1992.

"ISO CD 10303 - 42. Product Data Representation and Exchange - Part 42: Integrated Resources: Geometrical and Topological Representation," ISO Subcommittee 4, *Industrial Data and Global Manufacturing Programming Languages*. Secretariat: National Institute of Standards and Technology, Gaithersburg, MD, 1991.

Kruth, J.P., "Material Incress Manufacturing by Rapid Prototyping Techniques," *Annals of the CIRP* , vol. 40/2/, 1991.

Marcus, H.L., Beaman, J.J., Barlow,. J.W., Bourell, D.L., and Crawford, R.H., eds., "*Solid*

Rock, S. J., "Solid Freeform Fabrication and CAD System Interfacing," MS Thesis, (also RDRC Technical Report 91036) Rensselaer Design Research Center, Rensselaer Polytechnic Institute, Troy, NY, December 1991. [1991]

Prinz, F., Carnegie Mellon University, unpublished paper, 1988-1992.

Rock, S.J., and M.J. Wozny, "A Flexible File Format for Solid Freeform Fabrication," *Solid Freeform Fabrication Symposium 1991*, Marcus, Beaman, Barlow, Bourell, and Crawford, eds., Published by the University of Texas at Austin, August 1991. [1991a]

Rock, S.J., and M.J. Wozny, "Utilizing Topological Information to Increase Scan Vector Generation Efficiency," *Solid Freeform Fabrication Symposium 1991*, Marcus, Beaman, Barlow, Bourell, and Crawford, eds. , Published by The University of Texas at Austin, August 1991. [1991b]

Rock, S.J., and M.J. Wozny, "Generating Topological Information from a "Bucket of Facets"' *Solid Freeform Fabrication Symposium 1992*, Published by the University of Texas at Austin, August 1992.

Sankauratri, Sridhar. "Intelligent Scanning", unpublished preliminary version of M.S. Thesis. Rensselaer Polytechnic Institute, Troy, N.Y., August, 1992.

"Stereolithography Interface Specification," 3D Systems, Inc., June 1988.

Weiler, K., "Topological Structures for Geometric Modeling," Ph.D. Thesis, Rensselaer Polytechnic Institute, Troy, NY, August 1986. (Also RDRC Technical Report 86032)