

A Decomposition-Based Approach to Layered Manufacturing*

Ivaylo Ilinkin[†] Ravi Janardan^{†‡} Jayanth Majhi[§] Jörg Schwerdt^{¶||}
Michiel Smid^{§¶} Ram Sriram^{**}

September 20, 2001

Abstract

This paper introduces a new approach for improving the performance and versatility of *Layered Manufacturing* (LM), which is an emerging technology that makes it possible to build physical prototypes of 3D parts directly from their computer models using a “3D printer” attached to a personal computer.

Current LM processes work by viewing the computer model as a single, monolithic unit. By contrast, the approach proposed here decomposes the model into a small number of pieces, by intersecting it with a suitably chosen plane, builds each piece separately using LM, and then glues the pieces together to obtain the physical prototype. This approach allows large models to be built quickly in parallel. Furthermore, it is very efficient in its use of so-called support structures that are generated by the LM process.

This paper presents the first provably correct and efficient geometric algorithms to decompose polyhedral models so that the support requirements (support volume and area of contact) are minimized. Algorithms based on the plane-sweep paradigm are first given for convex polyhedra. These algorithms run in $O(n \log n)$ time for n -vertex convex polyhedra and work by generating expressions for the support volume and contact-area as a function of the height of the sweep plane, and optimizing them during the sweep. These algorithms are then generalized to non-convex polyhedra, which are considerably more difficult due to the complex structure of the supports. It is shown that, surprisingly, non-convex polyhedra can be handled by first identifying certain critical facets using a technique called cylindrical decomposition, and then applying the algorithm for convex polyhedra to these critical facets. The resulting algorithms run in

*A preliminary version of this paper will appear in the Proceedings of the Seventh International Workshop on Algorithms and Data Structures, Providence, RI, August 8–10, 2001.

[†]Department of Computer Science & Engineering, University of Minnesota, Minneapolis, MN 55455, U.S.A. {ilinkin, janardan}@cs.umn.edu. Research supported, in part, by National Science Foundation grant CCR-9712226 and by National Institute of Standards and Technology grant 60NANB8D0002. Commercial equipment and software, if any, are identified only in order to adequately specify certain procedures. In no case does such identification imply recommendation or endorsement by the National Institute of Standards and Technology, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

[‡]Portions of this work were done when RJ visited the University of Magdeburg, Germany under a joint grant from the National Science Foundation and Deutscher Akademischer Austauschdienst for international research.

[§]Synopsys Corporation, 700 East Middlefield Road, Mountain View, CA 94043. majhi@synopsys.com.

[¶]Fakultät für Informatik, Otto-von-Guericke-Universität Magdeburg, D-39106 Magdeburg, Germany. {schwerdt, michiel}@isg.cs.uni-magdeburg.de.

^{||}Portions of this work were done when MS and JS visited the University of Minnesota under a joint grant from Deutscher Akademischer Austauschdienst and the National Science Foundation for international research.

^{**}National Institute of Standards and Technology, Gaithersburg, MD 20899. sriram@cme.nist.gov.

$O(n^2 \log n)$ time. Also given is a method for controlling the size of the decomposition, so that the number of pieces generated is within a user-specified limit. Experimental results show that the proposed approach can achieve significant reduction in support requirements in both the convex and the non-convex case.

1 Introduction

The field of computer-aided design and manufacturing has now progressed to the point where an engineer can not only design and test a virtual model of a 3D part on a personal computer, but can also generate directly from the model a physical prototype of the part, using a relatively small and inexpensive “3D printer” attached to the computer [Jac92]. This technology, called *Rapid Prototyping* (RP), provides the engineer with an additional level of physical verification that makes it possible to detect and correct flaws that may, otherwise, have gone unnoticed. RP is used extensively in the automotive, aerospace, and medical industries.

At the heart of RP is a manufacturing technique called *Layered Manufacturing* (LM). The basic principle underlying LM is simple: the computer model is oriented suitably and sliced into thin layers by horizontal planes. The layers are then sent over a network to a fabrication device which “prints” them one by one, each on top of the previous one.

Figure 1 depicts a widely-used LM process called Stereolithography. In essence, the Stereolithography Apparatus (SLA) consists of a vat of light-sensitive liquid resin, a platform, and a laser. The input to the process (and to virtually all other LM processes) is a surface triangulation of the model in the industry-standard STL format. This format merely consists of an unordered list of the triangles, where each triangle is specified by listing its three vertices and its outward-directed unit-normal.¹ The model is oriented suitably, sliced by horizontal planes, and then built in the vertical direction as follows: Initially, the platform is below the surface of the resin at a depth equal to the layer thickness. The laser traces out the contour of the first slice on the surface and then fills in the interior in a raster-like pattern—a process called *hatching*. The resin then hardens to a depth equal to the slice thickness and forms the first layer, which rests on the platform. Next, the platform is lowered by an amount equal to the layer thickness, the vacated region is re-coated with resin, and the second layer is then built in the same way. Ideally, each layer after the first one should rest in its entirety on the previous one. In general, however, portions of a layer can overhang the previous layer, so additional structures, called *supports*, are needed to hold up the overhangs. Supports are generated automatically during the process itself. For this the model is analyzed *a priori* and a suitable description of the supports is generated and merged into the STL file. Once the part has been made, it is postprocessed to remove the supports and to improve the surface finish, which has a stair-stepped appearance owing to the discretization into layers.

1.1 Geometric issues in Layered Manufacturing

Unfortunately, the time requirements of present-day LM processes is quite high—often running into hours. Generally, an LM process consists of three phases: preprocessing, building, and postprocessing. Preprocessing includes repairing flaws in the STL file (e.g., gaps between facets, geometric singularities, etc. [Bai95, Bøh95]), deciding upon a suitable initial orientation for the model (or,

¹The simplicity of the STL format has led to its widespread adoption in the LM industry. However, the format does have some inadequacies and alternative formats have been proposed; see, for instance, [MKDP98, DKPS98].




Figure 1: *The Stereolithography apparatus.*

equivalently, the *build direction*), computing support requirements, generating and merging a description of the supports into the STL file, and slicing the model and supports. The building phase involves tracing and hatching each layer. Postprocessing includes removal of supports and improving part finish and accuracy.

All of these tasks are influenced by geometric considerations [Bai95, FF94]. In particular, the orientation of the part impacts the number of layers, the quantity of supports, the location of supports on the part, the extent to which supports “stick” to the part, and the surface finish and accuracy. The layer geometry affects the tool-path during hatching. The STL representation is not well-suited for the efficient computation of slices and supports as it does not contain any topological information about the model.

In current systems these issues are resolved in an *ad hoc* manner, through human intervention. For instance, the part is oriented by the operator, based on experience, so that the quantity of supports used is “small” and the finish is “good.” However, over the past few years the problem of automating these decisions via efficient and correct algorithms has been addressed by several researchers in computational geometry.

Asberg *et al.* [ABB⁺97] (see also [Bos95]), describe efficient algorithms to decide if a given model can be built without supports using Stereolithography. Majhi *et al.* [MJSG99] give algorithms to minimize the volume of supports and the area of contact between supports and the part for convex polyhedra, and show how to minimize the stair-step error for non-convex polyhedra. They also show how to reconcile multiple design criteria simultaneously [MJSS98] and give support optimization algorithms for non-convex polygons [MJS⁺99]. (See also [Maj98].) Schwerdt *et al.* show how to choose a build direction that protects prescribed facets from being damaged by supports [SSJ⁺00] and to minimize the number of layers [SSMJ99]. Recently, Agarwal and Desikan [AD00] have given an efficient algorithm to approximate a build direction which minimizes the contact-area for a convex polyhedron. Johnson [Joh99] shows how to compute support descriptions efficiently for a given build direction and McMains and Séquin [MS99] show how to slice a model efficiently using the plane sweep paradigm. Barequet and Kaplan [BK98] describe the design of a software front-end for automating many process planning functions in LM.

There has also been substantial effort from the computer-aided design community, which is primarily experimental/heuristic in nature. Representative work here includes that of Frank and Fadel [FF94] who consider support optimization in the framework of an expert system, Bablani and Bagchi [BB95] who were the first to quantify the notion of stair-step error, Allen and Dutta [AD95] who give heuristics for minimizing support contact-area, Bøhn [Bøh95] who addresses the problem of automatic repair of computer models, and Kulkarni and Dutta [KD96] and Dolenc and

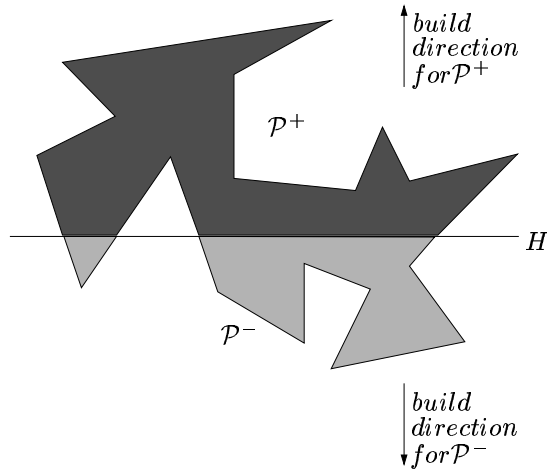


Figure 2: The decomposition-based approach, shown in 2D for convenience. The polyhedron \mathcal{P} is decomposed by a plane H into polyhedra \mathcal{P}^+ and \mathcal{P}^- , which are then built in the indicated directions.

Mäkelä [DM94] who consider the problem of slicing with variable layer thicknesses to capture fine detail in the model. The special issue by Stucki *et al.* [SBE95], contains a wealth of information on commercial and university-based work in RP and LM.

1.2 A decomposition-based approach

All the process-planning algorithms for LM that we are aware of work by viewing the model as a single, monolithic unit. By contrast, the approach that we propose works by decomposing the model into a small number of pieces, building each piece separately using LM, and then gluing the built pieces together to obtain the physical prototype. (The number of pieces generated can be controlled by the user.) Specifically, given a decomposition direction, we decompose the model by intersecting it with a suitable plane perpendicular to this direction; we then build the pieces that lie in the same halfspace in the direction given by the normal to the plane that is contained in the halfspace. Figure 2 illustrates this (in 2D, for convenience).

The decomposition-based approach has the advantage that it allows the construction of large models that cannot be accommodated in the workspace as a single piece. Moreover, the model can be built rapidly by building the pieces in parallel. Speed is particularly important in the “look and feel” prototyping that is the dominant use of RP today, since the goal here is to quickly get into the hands of the designer a physical version of the model to assess its general feel and appearance, and to rapidly iterate on the design until it is acceptable. In this context, the potential disadvantages of the decomposition-based approach (i.e., inaccuracies that may result from manually gluing the pieces together and reduced strength across the glued sections) are offset by the ability to build large models quickly in parallel.

A less obvious, but nevertheless crucial, advantage is that the support requirements of the decomposition-based approach are often much less than those of the conventional approach.² For

²They can never be more, since the conventional approach corresponds to the special case where the decomposition

instance, if a hollow sphere is built the conventional way, supports will be needed in the interior void and below the lower hemisphere; however, if it is built as two hemispheric shells, in opposite directions, then supports will be needed only in the regions previously occupied by the void, which results in an overall reduction in both support volume and contact-area. Our experiments confirm substantial savings for various other models as well. Reducing support requirements is important because this translates into lower material costs and faster build times.

1.3 Contributions

We give the first provably correct and efficient geometric algorithms to decompose polyhedral (i.e., STL) models, w.r.t. a given decomposition direction, so that the support contact-area and, independently, the support volume is minimized. We formalize this problem in Section 2. In Sections 3 and 4, we consider convex polyhedra and devise efficient, plane-sweep-based algorithms. We show how to generate expressions for the support volume and contact-area as a function of the height of the sweep plane, and how to optimize them during the sweep. We discuss our implementation of these algorithms and the results of our experiments on convex polyhedra with up to 200,000 vertices.

It may be argued (quite reasonably) that real-life models are generally not convex, so the direct relevance of the above algorithms is somewhat limited. Our original motivation for considering convex polyhedra was to get a handle on the problem for the non-convex case, which is considerably more difficult because of the complex structure of the supports (see Figure 3). Indeed, we will see that our approach for the convex case is one of the key building blocks to our solution for the non-convex case. In Section 5, we show how to handle non-convex polyhedra by first identifying certain critical facets (or parts thereof) using a technique called cylindrical decomposition [Mul93]. It turns out that the problem can then be solved by applying to these critical facets the algorithm developed for the convex case. We also give experimental results for typical non-convex polyhedra. Another key problem encountered in the non-convex case (but not in the convex case) is to keep small the number of pieces in the decomposition. We give an efficient technique for controlling the size of the decomposition of non-convex polyhedra, so that the number of pieces generated is within a user-specified limit.

To the best of our knowledge, the only prior work related to this approach is due to Fekete and Mitchell [FM97]. They consider the problem of decomposing a polyhedron into special polyhedra called histograms, which can be built on certain “base” facets with no supports. They prove that deciding if a polyhedron of genus zero can be decomposed into k histograms is NP-complete (this is true even in two dimensions for a polygon with holes). Our work differs from [FM97] in that we seek a decomposition (w.r.t. a given direction) into two polyhedra for which the total support requirement is minimized, but not necessarily zero.

2 Formalizing the problem

We denote by \mathcal{P} the polyhedron of interest. We assume the facets of \mathcal{P} are triangles and that its boundary is represented in some standard form, such as, for instance, a doubly-connected edge list [dBvKOS97] or a winged-edge structure [Bau75]. (If necessary, such a representation can be

plane coincides with the platform.




Figure 3: Support structures (in light shading), shown in 2D for convenience. Supports in the non-convex case (left) exhibit complexities not seen in the convex case (right): (i) They can rest partially on other parts of the polyhedron; (ii) only a fraction of a facet may be in contact with supports; (iii) parallel facets can also be in contact with supports.

computed easily from the standard STL representation of \mathcal{P} [MS99].) Let \mathbf{d} be a given *decomposition direction* (a unit-vector); we assume, w.l.o.g., that \mathbf{d} coincides with the positive z direction. Let H be any plane perpendicular to \mathbf{d} and intersecting \mathcal{P} ; we call H the *decomposition plane*. Let \mathcal{P}^+ be the closed polyhedron bounded by the facets of \mathcal{P} (or portions thereof) that are above H , and by the facet $\mathcal{P} \cap H$. (Note that \mathcal{P}^+ may consist of more than one connected component.) If no part of \mathcal{P} lies strictly above H , then \mathcal{P}^+ is taken to be empty. Define \mathcal{P}^- symmetrically w.r.t. the portion of \mathcal{P} that is below H . (Facet $\mathcal{P} \cap H$ appears in both \mathcal{P}^+ and \mathcal{P}^- , with outward normals $-\mathbf{d}$ and \mathbf{d} , respectively.) We define the *build direction* for \mathcal{P}^+ to be \mathbf{d} and for \mathcal{P}^- to be $-\mathbf{d}$, and we take H to be the *platform* for both polyhedra.

Let f be any facet of \mathcal{P} . We classify f , w.r.t. the given decomposition direction \mathbf{d} , as a *front facet*, a *back facet*, or a *parallel facet* of \mathcal{P} depending on whether the angle between the decomposition direction \mathbf{d} and the outward unit-normal, \mathbf{n}_f , of f is less than, greater than, or equal to 90° , respectively. (Note that f may be completely within \mathcal{P}^+ or \mathcal{P}^- , or may be partially in \mathcal{P}^+ and partially in \mathcal{P}^- . However, since the classification is done w.r.t. the fixed decomposition direction, rather than the build directions for \mathcal{P}^+ and \mathcal{P}^- , f has the same classification in \mathcal{P}^+ and/or \mathcal{P}^- as it does in \mathcal{P} .)

We now formalize the notion of supports. A facet of a polyhedron will need to be supported iff the angle between its outer normal and the build direction of the polyhedron is greater than 90° . This implies that the back facets of \mathcal{P}^+ and the front facets of \mathcal{P}^- will need to be supported. For concreteness, consider a back facet f of \mathcal{P}^+ . The *support polyhedron* for f is the closure of the set of all points $p \in \mathbb{R}^3$ such that p is not in the interior of \mathcal{P}^+ and the ray shot from p in direction \mathbf{d} first enters \mathcal{P}^+ through f . Informally, the support polyhedron of f is bounded from above by f , on the sides by vertical facets that “drop down” from the edges of f , and from below by the platform and/or portions of front facets of \mathcal{P}^+ . (If \mathcal{P}^+ is convex, then it is bounded from below by only the platform.) See Figure 3 for an example in 2D.

The *support contact-area* for \mathcal{P}^+ is the total surface area of \mathcal{P}^+ that is in contact with supports. It includes the area of all the back facets of \mathcal{P}^+ , except $\mathcal{P} \cap H$, and the areas of those portions of front facets and parallel facets that are in contact with supports. (Facet $\mathcal{P} \cap H$ rests on the platform and hence needs no supports. Note that while back facets are completely in contact with supports, front and parallel facets may be only partially in contact.) The *support volume* for \mathcal{P}^+

is the total volume of the support polyhedra for all back facets f of \mathcal{P}^+ (excluding $\mathcal{P} \cap H$). A symmetric discussion applies to the polyhedron \mathcal{P}^- .

We are now ready to state formally the problem that we wish to solve.

Problem 2.1 *Given a polyhedron \mathcal{P} , with n vertices, and a decomposition direction \mathbf{d} , compute a plane H perpendicular to \mathbf{d} which decomposes \mathcal{P} into polyhedra \mathcal{P}^+ and \mathcal{P}^- (as defined above) such that the total support contact-area or support volume is minimized when \mathcal{P}^+ and \mathcal{P}^- are built in directions \mathbf{d} and $-\mathbf{d}$, respectively. Additionally, if the user specifies an integer K , then the plane H should be optimal over all planes that generate no more than a total of K connected components of \mathcal{P}^+ and \mathcal{P}^- .*

3 Decomposing a convex polyhedron to minimize contact-area of supports

In this section and in Section 4, we assume that \mathcal{P} is convex. Thus \mathcal{P}^+ and \mathcal{P}^- are both convex and the support polyhedron for a back facet of \mathcal{P}^+ , or a front facet of \mathcal{P}^- , extends from the facet all the way to the platform, without intersecting any other facet (Figure 3). Furthermore, parallel facets will not be in contact with supports.³ These properties lead to an efficient algorithm for Problem 2.1, which we then incorporate into a solution for non-convex polyhedra.

Our approach is based on sweeping the plane H upwards starting from the xy -plane (we assume w.l.o.g. that \mathcal{P} lies above the xy -plane). Let H currently be at height h above the xy -plane and let f be any facet of \mathcal{P} . We classify f as an active or an inactive facet w.r.t. H depending on whether or not f is cut by H . Intuitively, an inactive facet is completely contained in \mathcal{P}^+ or \mathcal{P}^- and its contribution to the contact-area is not affected by small “local” movements of H . On the other hand, an active facet is contained partially in \mathcal{P}^+ and partially in \mathcal{P}^- , and small movements of H affect the facet’s contribution to the contact-area. Formally, we call f an *active facet w.r.t. H* if $H \cap f \neq \emptyset$ and at least one vertex of f is strictly above H . Otherwise, we call f an *inactive facet w.r.t. H* .

It follows that if f is an inactive front facet then its contribution to the total contact-area is $area(f)$ if it is in \mathcal{P}^- , and zero if it is in \mathcal{P}^+ , where $area(f)$ is the area of f . If f is an active front facet, then only the part, f^- , of f that lies below H is in \mathcal{P}^- , so f contributes $area(f^-)$ to the total contact-area. Symmetrically, if f is an inactive back facet, then its contribution to the total contact-area is $area(f)$ if it is in \mathcal{P}^+ , and zero if it is in \mathcal{P}^- . If f is an active back facet, then only the part f^+ of f that lies above H is in \mathcal{P}^+ , so f contributes $area(f^+)$ to the total contact-area.

The expression for the total contact-area of \mathcal{P}^+ and \mathcal{P}^- consists of two terms: the *inactive-area* term, which is the contact-area contributed by the inactive facets, and the *active-area* term, which is the contact-area contributed by the active facets. If we move H up or down, without crossing a vertex, then the inactive-area does not change, since inactive facets become active, and vice versa, only when H crosses a vertex. Therefore, the inactive-area term is simply a real number. However, the active-area changes because the fraction of an active facet that contributes to the total contact-area changes as H is moved. Lemma 3.1 below shows that the active-area term is a quadratic expression in h .

³Indeed, all references to facets in this section and in Section 4 mean non-parallel facets.




Figure 4: Intersection of H and active front facet f .

Lemma 3.1 *Let the plane H be at height h above the xy -plane. The total contact-area contributed by the active facets (i.e., the active-area) is of the form $Ah^2 + Bh + C$, where the coefficients A , B , and C depend only on the coordinates of the vertices of the active facets.*

Proof Let f be any active facet, with vertices $v_i = (x_i, y_i, z_i)$, $v_j = (x_j, y_j, z_j)$, and $v_k = (x_k, y_k, z_k)$. (Figure 4.) We will prove that the contact-area contributed by f is of the form $a_f h^2 + b_f h + c_f$, where the coefficients a_f , b_f , and c_f depend only on the coordinates of the vertices of f . This implies the result.

Note that by definition of an active facet, f is not contained in H . W.l.o.g., assume that f is a front facet, that $z_i < z_j \leq z_k$, and that H intersects f between v_i and v_j . (The other cases are similar and are discussed at the end of the proof.) Let H intersect edge $\overline{v_i v_j}$ at $v_{ij} = (x_{ij}, y_{ij}, z_{ij})$ and edge $\overline{v_i v_k}$ at $v_{ik} = (x_{ik}, y_{ik}, z_{ik})$. We compute these intersection points as follows.

Let ℓ_{ij} and ℓ_{ik} be the lines containing $\overline{v_i v_j}$ and $\overline{v_i v_k}$, respectively. For real-valued parameters t and r , the equations of these two lines are:

$$\ell_{ij} : x = x_i + (x_j - x_i)t, \quad y = y_i + (y_j - y_i)t, \quad z = z_i + (z_j - z_i)t,$$

and

$$\ell_{ik} : x = x_i + (x_k - x_i)r, \quad y = y_i + (y_k - y_i)r, \quad z = z_i + (z_k - z_i)r.$$

Since $z_{ij} = z_{ik} = h$, the parameters t and r are given by

$$t = (h - z_i)/(z_j - z_i), \quad r = (h - z_i)/(z_k - z_i).$$

Let $\alpha_j = (x_j - x_i)/(z_j - z_i)$, $\alpha_k = (x_k - x_i)/(z_k - z_i)$, $\beta_j = (y_j - y_i)/(z_j - z_i)$, and $\beta_k = (y_k - y_i)/(z_k - z_i)$. Then it follows that

$$x_{ij} = x_i + \alpha_j(h - z_i), \quad y_{ij} = y_i + \beta_j(h - z_i)$$

and

$$x_{ik} = x_i + \alpha_k(h - z_i), \quad y_{ik} = y_i + \beta_k(h - z_i).$$

Let f^- be the part of f below H ; f^- is a triangle with vertices v_i , v_{ij} , and v_{ik} . Then

$$2 \cdot \text{area}(f^-) = |(\mathbf{v}_{ij} - \mathbf{v}_i) \times (\mathbf{v}_{ik} - \mathbf{v}_i)|,$$

where $\mathbf{v}_{ij} - \mathbf{v}_i = \alpha_j(h - z_i)\mathbf{i} + \beta_j(h - z_i)\mathbf{j} + (h - z_i)\mathbf{k}$ and $\mathbf{v}_{ik} - \mathbf{v}_i = \alpha_k(h - z_i)\mathbf{i} + \beta_k(h - z_i)\mathbf{j} + (h - z_i)\mathbf{k}$, and $|\cdot|$ denotes the length of the associated vector. Therefore,

$$\begin{aligned} (\mathbf{v}_{ij} - \mathbf{v}_i) \times (\mathbf{v}_{ik} - \mathbf{v}_i) &= \begin{vmatrix} \mathbf{i} & \mathbf{j} & \mathbf{k} \\ \alpha_j(h - z_i) & \beta_j(h - z_i) & (h - z_i) \\ \alpha_k(h - z_i) & \beta_k(h - z_i) & (h - z_i) \end{vmatrix} \\ &= (h - z_i)^2(\beta_j - \beta_k)\mathbf{i} - (h - z_i)^2(\alpha_j - \alpha_k)\mathbf{j} + \\ &\quad (h - z_i)^2(\alpha_j\beta_k - \alpha_k\beta_j)\mathbf{k}. \end{aligned}$$

It follows that

$$\text{area}(f^-) = \frac{1}{2}(h - z_i)^2((\beta_j - \beta_k)^2 + (\alpha_j - \alpha_k)^2 + (\alpha_j\beta_k - \alpha_k\beta_j)^2)^{1/2}.$$

The coefficient of $(h - z_i)^2$ above is a constant which depends only on the coordinates of the vertices of f . In fact, it is easy to verify that this constant is equal to $\text{area}(f)/((z_j - z_i)(z_k - z_i))$. Hence the contribution of f to the active-area term is of the form $a_f h^2 + b_f h + c_f$, as claimed.

If $z_i = z_j$ or if H intersects f between v_j and v_k , then f^- is a quadrilateral and $\text{area}(f^-) = \text{area}(f) - \text{area}(f^+)$, where f^+ is the part of f above H . Since f^+ is a triangle, its area can be written as a quadratic expression similar to the one above. Moreover, $\text{area}(f)$ is a constant. It follows that $\text{area}(f^-)$ is of the claimed form. Finally, if f is a back facet, then a symmetric discussion applies.

Summing the area contributions of all active facets gives the lemma. ■

Remark 3.1 Note that the coefficient A in Lemma 3.1 is the sum of coefficients a_f for those active facets f that contribute triangular areas to the total contact-area minus the sum of coefficients a_g for those active facets g that contribute quadrilateral areas to the total contact-area. Therefore, it is possible that $A = 0$. Similarly, it is also possible that $B = 0$. In fact, the following example shows that A and B can be zero simultaneously.

Let \mathcal{P} be a cube which has been sheared by a small amount in the positive y direction. That is, if we view \mathcal{P} from the positive x -axis, then the top and bottom facets are parallel to the xy -plane, the front and back facets are parallel to the yz -plane, and the left and right facets are parallel but slanted rightwards. Let F and G be the left and right facets, respectively. Split F into triangles f and f' and G into triangles g and g' . Let f be the triangle of F with one vertex at the bottom and two at the top, and let g be the similar triangle of G . Note that $\text{area}(f) = \text{area}(g)$, and, w.r.t. the positive z decomposition direction, f is a front facet and g is a back facet. Let z_i , z_j , and z_k ($= z_j$) be the z -coordinates of the bottom, middle, and top vertices of f (and g).

Let H be anywhere between z_i and z_j . Then f 's contribution to the contact-area is a triangle whose area is $(h - z_i)^2 \text{area}(f) / ((z_j - z_i)(z_k - z_i))$, and g 's contribution is a quadrilateral whose area is $\text{area}(g) - (h - z_i)^2 \text{area}(g) / ((z_j - z_i)(z_k - z_i))$. Therefore, the total contribution of f and g is simply $\text{area}(g)$. A symmetric discussion shows that the total contribution of f' and g' is $\text{area}(f')$. It follows that the total contact-area of \mathcal{P} for this position of H is a constant equal to $\text{area}(g) + \text{area}(f')$, i.e., $A = B = 0$.

In the algorithm to follow, we will need to optimize the active-area formula over a certain range of h for which it is valid. We will show there how to handle these special cases.

3.1 The algorithm

We sweep the plane H upwards, starting at the lowest vertex of \mathcal{P} . We stop at each vertex and update the inactive-area and the active-area terms, based on the facets incident to the current vertex. We then minimize the active-area term (hence the total contact-area) as the position, h , of H ranges between the z -coordinates of the current vertex and the next one. This gives the optimal position for H between successive vertices. We repeat this for all vertices to find the globally-best position for H .

Here is the algorithm in more detail:

- In a preprocessing step, we sort the vertices of \mathcal{P} according to non-decreasing z -coordinates, as v_1, v_2, \dots, v_n (ties are broken arbitrarily). For each facet $f \in \mathcal{P}$, we determine whether it is a front, a back, or a parallel facet, and also compute $\text{area}(f)$. We set the active-area term identically equal to zero and the inactive-area term equal to the total area of the back facets. We set the current minimum contact-area equal to the sum of these two terms. We then scan the vertices in their sorted order.
- Let v_ℓ be the current vertex, $1 \leq \ell < n$. For each facet f incident to v_ℓ , we do the following:
 - Case 1:** v_ℓ is the lowest vertex of f . (Thus, f changes from inactive to active at v_ℓ .) If f is a back facet, then we subtract $\text{area}(f)$ from the inactive-area term. Using Lemma 3.1, we compute and add the expression $a_f h^2 + b_f h + c_f$ to the active-area term, thereby including the contribution of f^+ to the total contact-area when H is between z_ℓ and $z_{\ell+1}$. If f is a front facet, then we update only the active-area term to include the contribution of f^- to the total contact-area.
 - Case 2:** v_ℓ is the highest vertex of f . (Thus, f changes from active to inactive at v_ℓ .) This case is symmetric to Case 1. That is, if f is a back facet then we use Lemma 3.1 to compute and subtract the expression $a_f h^2 + b_f h + c_f$ from the active-area term, thereby removing the contribution of f^+ to the total contact-area when H is between z_ℓ and $z_{\ell+1}$. If f is a front facet, then we update the active area term to remove the contribution of f^- , and we add $\text{area}(f)$ to the inactive-area term.
 - Case 3:** v_ℓ is the middle vertex of f . Here f continues to be active, but the active-area term must be updated since H intersects a different edge of f above v_ℓ than it did below. We perform this update using Lemma 3.1.

After one of the three cases above is executed for all facets f incident to v_ℓ , we have a new active-area term $Ah^2 + Bh + C$, which is valid for h in the interval $[z_\ell, z_{\ell+1}]$. We can minimize

this term using standard techniques from calculus: If $A \neq 0$, then the minimum is attained at $h = -B/2A$ if $-B/2A \in [z_\ell, z_{\ell+1}]$, and at $h = z_\ell$ or at $h = z_{\ell+1}$ otherwise. If $A = 0$ and $B \neq 0$, then the minimum is attained at either $h = z_\ell$ or at $h = z_{\ell+1}$. If $A = B = 0$, then the active-area is constant in $[z_\ell, z_{\ell+1}]$, so we (arbitrarily) pick z_ℓ as the “optimum” value of h . In any case, once the optimum h has been determined, we compute the total contact-area for this h and update the current minimum contact-area, if necessary.

Theorem 3.1 *The contact-area version of Problem 2.1 can be solved in $O(n \log n)$ time for a convex polyhedron \mathcal{P} with n vertices ($O(n)$ time if the vertices are given in sorted order in the decomposition direction \mathbf{d}).*

Proof The correctness of the algorithm follows from the earlier discussion. The initialization time is dominated by the $O(n \log n)$ time to sort the vertices. At each vertex v_ℓ , the incident faces can be accessed in constant time apiece, since \mathcal{P} is given as a doubly-connected edge list. The time to update the active-area and inactive-area terms is proportional to the degree of v_ℓ . The minimization of the active-area term takes constant time. Summing over all vertices, the sweep time is $O(n)$ since the sum of the vertex degrees in \mathcal{P} is $O(n)$. The claimed time bound follows. ■

3.2 Experimental results

We have implemented the above algorithm and tested it on several convex polyhedra. The program is written in C++, and performs floating point computations in double-precision. The tests were done on a SUN Ultra 10 Sparc machine with 256 MB of main memory and a 440 MHz processor.

We generated two classes of test polyhedra. For the first class, we generated for each n , a set of n points at random on a sphere of radius 100 centered at the origin. (In our experiments, n ranged from 20,000 to 200,000, in steps of 20,000.) We then computed the convex hull of these n points using the `qhull` program [BDH96], and used it as our test polyhedron. (Since the n points were in convex position, all of them appeared on the convex hull. Also, we set the parameters to `qhull` such that it did not merge coplanar facets; thus all facets were triangles.) Table 1 shows the running times of our program and the minimum contact-area computed. To illustrate the savings realized by decomposition, we also computed the contact-area without decomposition in each case.

As one might expect, the test polyhedra generated above were quite symmetric about the origin. As a result, the optimum decomposition plane was always close to the xy -plane. Therefore, we also generated a second set of polyhedra that did not exhibit such symmetry. For this, we used `qhull` to generate $n + 1$ points at random on a cone whose major axis was along the z -axis, for n in the range 20,000 to 200,000 (the additional point was the apex of the cone). We then rotated the cone by a randomly chosen angle to make it non-symmetric about the origin. For each n , we generated two non-symmetric test polyhedra in this fashion. Table 2 shows our results; here “angle” refers to the angle by which the cone was rotated. For each input size, the running times on the two polyhedra were nearly the same; therefore, we averaged the times. This closeness in run times is to be expected since the complexity of the algorithm depends primarily on the graph structure of \mathcal{P} , which is the same regardless of orientation. In fact, since the graph is fully triangulated, the number of edges and faces are the same for different polyhedra with the same number of vertices. This is borne out by Tables 1 and 2.

Support Contact-area				
#verts n	min contact-area	h_{min}	run time (s)	non-decomp. contact-area
20,000	579.0	-.14	.6	62,708.1
40,000	424.2	.04	1.2	62,832.6
60,000	338.6	-.04	2.0	62,803.7
80,000	286.6	-.05	2.7	62,802.5
100,000	274.6	.02	3.5	62,840.0
120,000	243.6	.05	4.2	62,864.0
140,000	226.4	0	5.0	62,824.7
160,000	211.6	-.01	5.8	62,833.2
180,000	191.2	.03	6.5	62,845.1
200,000	174.7	-.02	7.4	62,814.9

Table 1: Minimum support contact-area for convex polyhedra generated from random points on a sphere of radius 100. Here “non-decomp. contact-area” refers to the support contact-area when the polyhedra are built without decomposition; observe the significant reductions achieved via decomposition. Note that the support contact-area without decomposition is roughly equal to the area of the lower hemisphere, as it should be.

3.3 Discussion

One might wonder if there is always an optimal solution where the plane H passes through a vertex of \mathcal{P} , thus obviating the need to minimize the active-area formula in between successive vertices. In this section, we show that this is not the case by proving that for the pyramid \mathcal{P} of Figure 5 the optimal plane does not contain any vertex.

In what follows, we will use the notation $\langle v_i, v_j, v_k \rangle$ to denote a facet (triangle) of \mathcal{P} that is bounded by vertices v_i, v_j , and v_k .

It is easy to verify that

$$2 \cdot \text{area}(\langle v_1, v_3, v_2 \rangle) = |(\mathbf{v}_3 - \mathbf{v}_1) \times (\mathbf{v}_2 - \mathbf{v}_1)| = \sqrt{6},$$

$$2 \cdot \text{area}(\langle v_1, v_2, v_4 \rangle) = |(\mathbf{v}_4 - \mathbf{v}_1) \times (\mathbf{v}_2 - \mathbf{v}_1)| = \sqrt{6},$$

$$2 \cdot \text{area}(\langle v_2, v_3, v_4 \rangle) = |(\mathbf{v}_3 - \mathbf{v}_2) \times (\mathbf{v}_4 - \mathbf{v}_2)| = \sqrt{8}.$$

Note that facet $\langle v_1, v_3, v_4 \rangle$ is a parallel facet for the indicated decomposition direction \mathbf{d} , and, therefore, does not contribute to the contact-area.

We now compute the total contact-area for \mathcal{P}^+ and \mathcal{P}^- for each decomposition plane H that passes through a vertex of \mathcal{P} .

Case 1: H passes through v_1 , i.e., H is the plane $z = -1$.

In this case, \mathcal{P}^- is empty and the (active) back facets $\langle v_1, v_3, v_2 \rangle$ and $\langle v_1, v_2, v_4 \rangle$ contribute $\text{area}(\langle v_1, v_3, v_2 \rangle)$ and $\text{area}(\langle v_1, v_2, v_4 \rangle)$, respectively, to the contact-area for \mathcal{P}^+ . Therefore, the total contact-area is $\sqrt{6}/2 + \sqrt{6}/2 = 2.449$.

Support Contact-area					
#verts n	angle	min contact-area	h_{min}	mean run time (s)	non-decomp. contact-area
20,001	200	7,752.6	-3.2	.6	55,300.5
	37	11,705.0	2.3		57,111.7
40,001	40	12,193.1	2.7	1.3	58,304.3
	75	3,973.7	.7		52,005.6
60,001	240	7,037.8	-.5	2.0	55,082.1
	112	5,315.5	-.5		55,068.2
80,001	80	2,714.5	0	2.8	52,272.1
	150	10,108.5	-1.5		51,825.6
10,0001	280	2,649.1	0	3.6	52,160.8
	187	3,447.8	-3.6		59,820.0
120,001	120	7,018.4	-.2	4.3	55,068.5
	225	11,908.1	-1.0		49,168.8
140,001	320	12,209.7	1.2	5.1	58,237.1
	263	1,967.0	.1		54,011.3
160,001	160	7,657.2	-2.1	5.9	55,325.7
	300	6,962.6	.2		51,425.1
180,001	0	5.0	1.7	6.7	44,798.2
	338	8,175.7	1.5		51,938.3
200,001	200	7,625.8	-2.1	7.5	55,297.5
	15	6,198.2	3.4		49,509.5

Table 2: Minimum support contact-area for convex polyhedra generated from random points on a cone; each polyhedron has been rotated by the indicated angle to make it non-symmetric about the origin. Here “non-decomp. contact-area” refers to the support contact-area when the polyhedra are built without decomposition; observe the significant reductions achieved via decomposition.




Figure 5: A polyhedron for which the decomposition plane realizing the minimum contact-area does not pass through a vertex.

Case 2: H passes through v_2 , i.e., H is the plane $z = 0$.

Let H intersect edges $\overline{v_1v_3}$ and $\overline{v_1v_4}$ at points v_{13} and v_{14} , respectively. We have $v_{13} = (0.5, 0, 0)$ and $v_{14} = (-0.5, 0, 0)$. There are no facets that contribute to the contact-area for \mathcal{P}^- . The (active) back facets $\langle v_1, v_3, v_2 \rangle$ and $\langle v_1, v_2, v_4 \rangle$ contribute $area(\langle v_2, v_{13}, v_3 \rangle)$ and $area(\langle v_2, v_{14}, v_4 \rangle)$, respectively, to the contact-area for \mathcal{P}^+ . We have

$$2 \cdot area(\langle v_2, v_{13}, v_3 \rangle) = |(\mathbf{v}_{13} - \mathbf{v}_2) \times (\mathbf{v}_3 - \mathbf{v}_2)| = \sqrt{1.5},$$

$$2 \cdot area(\langle v_2, v_{14}, v_4 \rangle) = |(\mathbf{v}_{14} - \mathbf{v}_2) \times (\mathbf{v}_4 - \mathbf{v}_2)| = \sqrt{1.5}.$$

Therefore, the total contact-area is $\sqrt{1.5}/2 + \sqrt{1.5}/2 = 1.225$.

Case 3: H passes through v_3 (and v_4), i.e., H is the plane $z = 1$.

Here \mathcal{P}^+ is empty. Only the (inactive) front facet $\langle v_2, v_3, v_4 \rangle$ contributes to the contact-area for \mathcal{P}^- . Its contribution is $area(\langle v_2, v_3, v_4 \rangle) = \sqrt{2} = 1.414$, which is also the total contact-area.

Now, let H be the plane $z = 0.5$. This plane does not pass through any vertex of \mathcal{P} . We will show that the total contact-area for this plane is smaller than any of the contact-areas computed in Cases 1–3 above.

Let H intersect edges $\overline{v_1v_3}$, $\overline{v_1v_4}$, $\overline{v_2v_3}$, and $\overline{v_2v_4}$ at points v_{13} , v_{14} , v_{23} , and v_{24} , respectively. We have $v_{13} = (0.75, 0, 0.5)$, $v_{14} = (-0.75, 0, 0.5)$, $v_{23} = (0.5, -0.5, 0.5)$, and $v_{24} = (-0.5, -0.5, 0.5)$.

Facet $\langle v_2, v_3, v_4 \rangle$ is a (active) front facet and contributes $area(\langle v_2, v_{23}, v_{24} \rangle)$ to the contact-area for \mathcal{P}^- . Facets $\langle v_1, v_3, v_2 \rangle$ and $\langle v_1, v_2, v_4 \rangle$ are (active) back facets, and contribute $area(\langle v_{23}, v_{13}, v_3 \rangle)$ and $area(\langle v_{14}, v_4, v_{24} \rangle)$, respectively, to the contact-area for \mathcal{P}^+ . We have

$$2 \cdot \text{area}(\langle v_2, v_{23}, v_{24} \rangle) = |(\mathbf{v}_{23} - \mathbf{v}_2) \times (\mathbf{v}_{24} - \mathbf{v}_2)| = 1/\sqrt{2},$$

$$2 \cdot \text{area}(\langle v_{23}, v_{13}, v_3 \rangle) = |(\mathbf{v}_{23} - \mathbf{v}_3) \times (\mathbf{v}_{13} - \mathbf{v}_3)| = \sqrt{1.5}/4,$$

$$2 \cdot \text{area}(\langle v_{14}, v_4, v_{24} \rangle) = |(\mathbf{v}_{14} - \mathbf{v}_4) \times (\mathbf{v}_{24} - \mathbf{v}_4)| = \sqrt{1.5}/4.$$

Thus, the total contact-area is $(1/2)1/\sqrt{2} + (1/2)\sqrt{1.5}/4 + (1/2)\sqrt{1.5}/4 = 0.660$, which is less than any of the contact-areas computed in Cases 1–3. This proves that the decomposition plane that realizes the minimum contact-area need not pass through any vertex of \mathcal{P} . It turns out that, for this example, the optimal plane is $H : z = 0.46$, and the corresponding total contact-area is 0.656. (This was computed by our program in Section 3.2.)

4 Decomposing a convex polyhedron to minimize volume of supports

In this section, we describe an algorithm to decompose \mathcal{P} such that when the resulting polyhedra \mathcal{P}^+ and \mathcal{P}^- are built in directions \mathbf{d} and $-\mathbf{d}$, respectively, the total volume of supports is minimized. Each support polyhedron for \mathcal{P}^+ is bounded from above by a back facet, on the sides by vertical edges emanating from the edges of the facet, and from below by the decomposition plane H . A symmetric discussion applies for \mathcal{P}^- .

For any given decomposition plane $H : z = h$, the expression for the support volume consists of two terms: The *inactive-volume* term and the *active-volume* term, which are, respectively, the total volumes of the support polyhedra contributed by the inactive and the active facets. If we move H up or down without crossing a vertex, the active-volume as well as the inactive-volume changes (unlike contact-area, where only the active-area changes). Lemma 4.1 below proves that the active-volume term is cubic in h , and Lemma 4.2 shows that the inactive-volume is linear in h .

Lemma 4.1 *Let the plane H be at height h above the xy -plane. The total support volume contributed by the active facets (i.e., the active-volume) is of the form $Ah^3 + Bh^2 + Ch + D$, where the coefficients A , B , C , and D depend only on the coordinates of the vertices of the active facets.*

Proof Let f be any active facet, with vertices $v_i = (x_i, y_i, z_i)$, $v_j = (x_j, y_j, z_j)$, and $v_k = (x_k, y_k, z_k)$, where $z_i < z_j$, $z_i < z_k$ and let the vertices v_i , v_j , and v_k be in counterclockwise order when viewed from the outside of \mathcal{P} (i.e., $(\mathbf{v}_j - \mathbf{v}_i) \times (\mathbf{v}_k - \mathbf{v}_i)$ has the same direction as the outward unit-normal, \mathbf{n}_f , to f). Clearly, it suffices to prove that the support volume contributed by f is of the form $a_f h^3 + b_f h^2 + c_f h + d_f$, where the coefficients a_f , b_f , c_f , and d_f depend only on the coordinates of the vertices of f .

Note that, by definition of an active facet, f is not contained in H . Assume that H intersects edges $\overline{v_i v_j}$ and $\overline{v_i v_k}$, so that f^- is a triangle and f^+ is a quadrilateral. (The other cases are similar and are discussed at the end of the proof.) If f is a front facet, then it contributes a support polyhedron only to \mathcal{P}^- , whereas if f is a back facet, then it contributes a support polyhedron only to \mathcal{P}^+ . Let us denote this support polyhedron by R_f . Let v'_i , v'_j , and v'_k be the vertical projections




Figure 6: The support polyhedron, R_f , contributed by an active front facet f . R_f is defined by the vertices v_i , v_i' , v_{ij} , and v_{ik} .

of v_i , v_j , and v_k , respectively, on H . If f is a front facet, then R_f is defined by the vertices v_i , v_i' , v_{ij} , and v_{ik} (Figure 6); if f is a back facet, then R_f is defined by the vertices v_{ij} , v_{ik} , v_j' , v_k' , v_j , and v_k (Figure 7). Our goal is to find a formula for the volume, $vol(R_f)$, of R_f in each case.

Let S_j , $1 \leq j \leq r$ be the facets of R_f , for some positive integer r . Let \mathbf{N}_j be the outward unit-normal to S_j ("outward" w.r.t. R_f , not \mathcal{P}). Let Q_j be any point on S_j and let \mathbf{Q}_j be the position vector of Q_j . From [Gol91] we have

$$vol(R_f) = \frac{1}{3} \left| \sum_j (\mathbf{Q}_j \cdot \mathbf{N}_j) area(S_j) \right|. \quad (1)$$

We next consider the two possibilities for f separately.

Case 1: f is a front facet of \mathcal{P} (Figure 6).

In this case R_f is bounded by the following facets:

$$\begin{aligned} S_1 &: \langle v_i, v_{ij}, v_i' \rangle \\ S_2 = f^- &: \langle v_i, v_{ik}, v_{ij} \rangle \\ S_3 &: \langle v_i, v_i', v_{ik} \rangle \\ S_4 &: \langle v_i', v_{ij}, v_{ik} \rangle \end{aligned}$$

The following table shows \mathbf{N}_j and the point Q_j for each facet S_j . We have chosen Q_j such that subsequent calculations become simple.

Facet	Point Q_j	Normal \mathbf{N}_j
S_1	v_i	$(\mathbf{v}_j - \mathbf{v}_i) \times \mathbf{d} / (\mathbf{v}_j - \mathbf{v}_i) \times \mathbf{d} $
S_2	v_i	$-\mathbf{n}_f$
S_3	v_i	$\mathbf{d} \times (\mathbf{v}_k - \mathbf{v}_i) / \mathbf{d} \times (\mathbf{v}_k - \mathbf{v}_i) $
S_4	v'_i	\mathbf{d}

Note that all the dot products $\mathbf{Q}_j \cdot \mathbf{N}_j$, except for $\mathbf{Q}_4 \cdot \mathbf{N}_4$, depend only on f (and \mathbf{d} , which is fixed); the dot product $\mathbf{Q}_4 \cdot \mathbf{N}_4 = h$. We now generate an expression for each $area(S_j)$.

Expression for $area(S_1)$:

Consider the vertical projection, v''_i , of v_i on the plane passing through v_j and parallel to H .⁴ Facet S_1 is simply the triangular portion of $\langle v_i, v_j, v''_i \rangle$ lying below H . From the proof of Lemma 3.1

$$area(S_1) = (h - z_i)^2 \cdot area(\langle v_i, v_j, v''_i \rangle) / ((z_j - z_i)(z_j - z_i)),$$

where $area(\langle v_i, v_j, v''_i \rangle) = |(\mathbf{v}_j - \mathbf{v}_i) \times (\mathbf{v}''_i - \mathbf{v}_i)|/2$, and $\mathbf{v}''_i - \mathbf{v}_i$ has components $(0, 0, z_j - z_i)$. It follows that $area(S_1)$ can be written in the form

$$area(S_1) = a_f^{(1)} h^2 + b_f^{(1)} h + c_f^{(1)}, \quad (2)$$

for some coefficients $a_f^{(1)}$, $b_f^{(1)}$, and $c_f^{(1)}$ that depend only on the coordinates of the vertices of f .

Expression for $area(S_2)$:

From the proof of Lemma 3.1 it is immediate that

$$area(S_2) = a_f^{(2)} h^2 + b_f^{(2)} h + c_f^{(2)}, \quad (3)$$

for some coefficients $a_f^{(2)}$, $b_f^{(2)}$, and $c_f^{(2)}$ that depend only on the coordinates of the vertices of f .

Expression for $area(S_3)$:

The discussion here is identical to the one for $area(S_1)$, except that we use the vertical projection, v'''_i , of v_i on the plane passing through v_k and parallel to H . Therefore,

$$area(S_3) = (h - z_i)^2 \cdot area(\langle v_i, v_k, v'''_i \rangle) / ((z_k - z_i)(z_k - z_i)),$$

where $area(\langle v_i, v_k, v'''_i \rangle) = |(\mathbf{v}_k - \mathbf{v}_i) \times (\mathbf{v}'''_i - \mathbf{v}_i)|/2$, and $\mathbf{v}'''_i - \mathbf{v}_i = (0, 0, z_k - z_i)$. It follows that $area(S_3)$ can be written in the form

$$area(S_3) = a_f^{(3)} h^2 + b_f^{(3)} h + c_f^{(3)}, \quad (4)$$

⁴To avoid clutter, we do not show v''_i in Figure 6.

for some coefficients $a_f^{(3)}$, $b_f^{(3)}$, and $c_f^{(3)}$ that depend only on the coordinates of the vertices of f .

Expression for $area(S_4)$:

Now, $area(S_4)$ is simply the area of the projection of f^- on H , and is hence equal to $area(f^-)$ times the cosine of the angle between H and the plane containing f^- . Therefore, $area(S_4) = area(f^-)(\mathbf{n}_f \cdot \mathbf{d})$.

Recall from the proof of Lemma 3.1 that $area(f^-)$ is quadratic in h . It follows that $area(S_4)$ can be written in the form

$$area(S_4) = a_f^{(4)}h^2 + b_f^{(4)}h + c_f^{(4)}, \quad (5)$$

for some coefficients $a_f^{(4)}$, $b_f^{(4)}$, and $c_f^{(4)}$ that depend only on the coordinates of the vertices of f .

Let us denote by $active-volume_f(h)$ the support volume contributed by an active facet f w.r.t. the plane $H : z = h$. From Equations (2)–(5) and the fact that $\mathbf{Q}_4 \cdot \mathbf{N}_4 = h$, we have

$$\begin{aligned} active-volume_f(h) &= \\ &= \frac{1}{3} \left| \sum_{j=1}^4 (\mathbf{Q}_j \cdot \mathbf{N}_j) (a_f^{(j)} h^2 + b_f^{(j)} h + c_f^{(j)}) \right| \\ &= \frac{1}{3} \left| \sum_{j=1}^3 (\mathbf{Q}_j \cdot \mathbf{N}_j) (a_f^{(j)} h^2 + b_f^{(j)} h + c_f^{(j)}) + h (a_f^{(4)} h^2 + b_f^{(4)} h + c_f^{(4)}) \right| \\ &= \frac{1}{3} \left| a_f^{(4)} h^3 + \left(b_f^{(4)} + \sum_{j=1}^3 (\mathbf{Q}_j \cdot \mathbf{N}_j) a_f^{(j)} \right) h^2 + \right. \\ &\quad \left. \left(c_f^{(4)} + \sum_{j=1}^3 (\mathbf{Q}_j \cdot \mathbf{N}_j) b_f^{(j)} \right) h + \sum_{j=1}^3 (\mathbf{Q}_j \cdot \mathbf{N}_j) c_f^{(j)} \right|. \end{aligned}$$

Therefore,

$$active-volume_f(h) = a_f h^3 + b_f h^2 + c_f h + d_f, \quad (6)$$

for some coefficients a_f , b_f , c_f , and d_f that depend only on the coordinates of the vertices of f .

Case 2: f is a back facet of \mathcal{P} (Figure 7).




Figure 7: The support polyhedron, R_f , contributed by an active back facet f . R_f is defined by the vertices $v_j, v_k, v_{ij}, v_{ik}, v'_j,$ and v'_k .

In this case R_f is bounded by the following facets:

$$\begin{aligned}
 S_1 &: \langle v_{ik}, v'_k, v_k \rangle \\
 S_2 = f^+ &: \langle v_{ik}, v_k, v_j, v_{ij} \rangle \\
 S_3 &: \langle v_{ij}, v_j, v'_j \rangle \\
 S_4 &: \langle v_j, v_k, v'_k, v'_j \rangle \\
 S_5 &: \langle v_{ij}, v'_j, v'_k, v_{ik} \rangle
 \end{aligned}$$

The following table shows \mathbf{N}_j and the point Q_j for each facet S_j . Again, we have chosen Q_j such that subsequent calculations become simple.

Facet	Point Q_j	Normal \mathbf{N}_j
S_1	v_k	$\mathbf{d} \times (\mathbf{v}_i - \mathbf{v}_k) / \mathbf{d} \times (\mathbf{v}_i - \mathbf{v}_k) $
S_2	v_k	$-\mathbf{n}_f$
S_3	v_j	$\mathbf{d} \times (\mathbf{v}_j - \mathbf{v}_i) / \mathbf{d} \times (\mathbf{v}_j - \mathbf{v}_i) $
S_4	v_j	$\mathbf{d} \times (\mathbf{v}_k - \mathbf{v}_j) / \mathbf{d} \times (\mathbf{v}_k - \mathbf{v}_j) $
S_5	v'_k	$-\mathbf{d}$

Note that all the dot products $\mathbf{Q}_j \cdot \mathbf{N}_j$, except for $\mathbf{Q}_5 \cdot \mathbf{N}_5$, depend only on f (and \mathbf{d} , which is fixed); the dot product $\mathbf{Q}_5 \cdot \mathbf{N}_5 = -h$. We now generate an expression for each $area(S_j)$.

Expression for $area(S_1)$:

Consider the vertical projection, v_k'' , of v_k on the plane passing through v_i and parallel to H .⁵ Facet S_1 is simply the triangular portion of $\langle v_i, v_k, v_k'' \rangle$ lying above H . From the proof of Lemma 3.1

$$\text{area}(S_1) = (h - z_k)^2 \cdot \text{area}(\langle v_i, v_k, v_k'' \rangle) / ((z_i - z_k)(z_i - z_k)),$$

where $\text{area}(\langle v_i, v_k, v_k'' \rangle) = |(\mathbf{v}_i - \mathbf{v}_k) \times (\mathbf{v}_k'' - \mathbf{v}_k)|/2$, and $\mathbf{v}_k'' - \mathbf{v}_k$ has components $(0, 0, z_i - z_k)$. It follows that $\text{area}(S_1)$ can be written in the form

$$\text{area}(S_1) = a_f^{(1)} h^2 + b_f^{(1)} h + c_f^{(1)}, \quad (7)$$

for some coefficients $a_f^{(1)}$, $b_f^{(1)}$, and $c_f^{(1)}$ that depend only on the coordinates of the vertices of f .⁶

Expression for $\text{area}(S_2)$:

From the proof of Lemma 3.1 it is immediate that

$$\text{area}(S_2) = a_f^{(2)} h^2 + b_f^{(2)} h + c_f^{(2)}, \quad (8)$$

for some coefficients $a_f^{(2)}$, $b_f^{(2)}$, and $c_f^{(2)}$ that depend only on the coordinates of the vertices of f .

Expression for $\text{area}(S_3)$:

The discussion is identical to the one above for $\text{area}(S_1)$, except that we use the projection v_j'' of v_j on the plane passing through v_i and parallel to H . Therefore,

$$\text{area}(S_3) = (h - z_j)^2 \cdot \text{area}(\langle v_i, v_j, v_j'' \rangle) / ((z_i - z_j)(z_i - z_j)),$$

where $\text{area}(\langle v_i, v_j, v_j'' \rangle) = |(\mathbf{v}_i - \mathbf{v}_j) \times (\mathbf{v}_j'' - \mathbf{v}_j)|/2$, and $\mathbf{v}_j'' - \mathbf{v}_j$ has components $(0, 0, z_i - z_j)$. It follows that $\text{area}(S_3)$ can be written in the form

$$\text{area}(S_3) = a_f^{(3)} h^2 + b_f^{(3)} h + c_f^{(3)}, \quad (9)$$

for some coefficients $a_f^{(3)}$, $b_f^{(3)}$, and $c_f^{(3)}$ that depend only on the coordinates of the vertices of f .

Expression for $\text{area}(S_4)$:

Facet S_4 is a trapezoid with parallel sides $\overline{v_k v_k'}$ and $\overline{v_j v_j'}$, of lengths $z_k - h$ and $z_j - h$, respectively, and height $|\overline{v_j' v_k'}| = ((x_j - x_k)^2 + (y_j - y_k)^2)^{1/2}$. Therefore,

⁵Again, to avoid clutter, we do not show v_k'' in Figure 7.

⁶These coefficients are, of course, different from the ones we use in Case 1.

$$\begin{aligned}
\text{area}(S_4) &= (|\overline{v_k v'_k}| + |\overline{v_j v'_j}|) \cdot |\overline{v'_j v'_k}|/2 \\
&= -|\overline{v'_j v'_k}|h + |\overline{v'_j v'_k}|(z_j + z_k)/2
\end{aligned}$$

It follows that $\text{area}(S_4)$ can be written in the form

$$\text{area}(S_4) = a_f^{(4)}h^2 + b_f^{(4)}h + c_f^{(4)}, \quad (10)$$

where $a_f^{(4)} = 0$, and $b_f^{(4)}$ and $c_f^{(4)}$ depend only on the coordinates of the vertices of f .

Expression for $\text{area}(S_5)$:

Let f_H be the vertical projection of f on the plane H , with vertices v'_i , v'_j , and v'_k , corresponding to v_i , v_j , and v_k , respectively. (Note that v_{ij} and v_{ik} get projected to themselves.) Then

$$\text{area}(S_5) = \text{area}(f_H) - \text{area}(\langle v'_i, v_{ij}, v_{ik} \rangle).$$

Now $\text{area}(\langle v'_i, v_{ij}, v_{ik} \rangle)$ has already been computed in Case 1 (Equation (5)). Moreover, $\text{area}(f_H) = ((\mathbf{v}_j - \mathbf{v}_i) \times (\mathbf{v}_k - \mathbf{v}_i)) \cdot \mathbf{d})/2$.

Since the former is quadratic in h and the latter is a constant, it follows that $\text{area}(S_5)$ is quadratic in h . Thus,

$$\text{area}(S_5) = a_f^{(5)}h^2 + b_f^{(5)}h + c_f^{(5)}, \quad (11)$$

for some coefficients $a_f^{(5)}$, $b_f^{(5)}$, and $c_f^{(5)}$ that depend only on the coordinates of the vertices of f .

From Equations (7)–(11) and the fact that $\mathbf{Q}_5 \cdot \mathbf{N}_5 = -h$, we have

$$\begin{aligned}
\text{active-volume}_f(h) &= \\
&= \frac{1}{3} \left| \sum_{j=1}^5 (\mathbf{Q}_j \cdot \mathbf{N}_j) (a_f^{(j)}h^2 + b_f^{(j)}h + c_f^{(j)}) \right| \\
&= \frac{1}{3} \left| \sum_{j=1}^4 (\mathbf{Q}_j \cdot \mathbf{N}_j) (a_f^{(j)}h^2 + b_f^{(j)}h + c_f^{(j)}) - h(a_f^{(5)}h^2 + b_f^{(5)}h + c_f^{(5)}) \right| \\
&= \frac{1}{3} \left| -a_f^{(5)}h^3 + \left(-b_f^{(5)} + \sum_{j=1}^4 (\mathbf{Q}_j \cdot \mathbf{N}_j) a_f^{(j)} \right) h^2 + \right. \\
&\quad \left. \left(-c_f^{(5)} + \sum_{j=1}^4 (\mathbf{Q}_j \cdot \mathbf{N}_j) b_f^{(j)} \right) h + \sum_{j=1}^4 (\mathbf{Q}_j \cdot \mathbf{N}_j) c_f^{(j)} \right|.
\end{aligned}$$

Therefore,

$$\text{active-volume}_f(h) = a_f h^3 + b_f h^2 + c_f h + d_f, \quad (12)$$

for some coefficients a_f , b_f , c_f , and d_f that depend only on the coordinates of the vertices of f .

The lemma now follows if we sum the active-volume contributions of all front and back facets, using Equation (6) and Equation (12), respectively.

We now discuss how to handle the other cases mentioned at the beginning of the proof. Let H intersect f such that f^- is a quadrilateral and f^+ is a triangle. If f is a back facet, then we can use the discussion from Case 1 to compute the active-volume. Suppose that $z_i > z_j$, $z_i > z_k$, and the vertices v_i , v_j , and v_k are in counterclockwise order with respect to \mathbf{n}_f . All the formulas for the facet areas still apply. However, the directions of the normals \mathbf{N}_j must now be reversed (except for \mathbf{N}_2). Also, now $\mathbf{Q}_4 \cdot \mathbf{N}_4 = -h$, and therefore the coefficients of $\text{area}(S_4)$ in the active-volume formula must be taken with opposite signs.

If f is a front facet we use the discussion from Case 2 to compute the active-volume. Again, suppose that $z_i > z_j$, $z_i > z_k$, and the vertices v_i , v_j , and v_k are given in counterclockwise order w.r.t. \mathbf{n}_f . All the formulas for the facet areas still apply. However, the directions of the normals \mathbf{N}_j must now be reversed (except for \mathbf{N}_2). Also, now $\mathbf{Q}_5 \cdot \mathbf{N}_5 = h$, and therefore the coefficients of $\text{area}(S_5)$ in the active-volume formula must be taken with opposite signs. ■

Next we obtain an expression for the inactive-volume.

Lemma 4.2 *Let the plane H be at height h above the xy -plane. The total support volume contributed by the inactive facets (i.e., the inactive-volume) is of the form $Ch + D$, where the coefficients C and D depend only on the coordinates of the vertices of the inactive facets.*

Proof Let f be any inactive back facet with vertices $v_i = (x_i, y_i, z_i)$, $v_j = (x_j, y_j, z_j)$, and $v_k = (x_k, y_k, z_k)$, where $z_i \leq z_j \leq z_k$; see Figure 8. (The other cases are similar and discussed at the end of the proof.) We will prove that the support volume contributed by f is of the form $c_f h + d_f$ where the coefficients c_f and d_f depend only on the coordinates of the vertices of f . This implies the result.

Since f is an inactive back facet it will contribute to the support volume only if H is below f . Let v_j'' and v_k'' be the projections of v_j and v_k , respectively, on the plane passing through v_i and parallel to H . Let v_i' , v_j' , and v_k' be the projections of v_i , v_j , v_k , respectively, on H . We would like to find a formula for the volume of the polyhedron, T_f , defined by the vertices v_i , v_j , v_k , v_i' , v_j' and v_k' .

T_f consists of two polyhedra. The first, denoted T_f^{const} , is defined by v_i , v_j , v_k , v_j'' , and v_k'' , and the second, denoted T_f^{var} , is defined by v_i , v_i' , v_j' , v_k' , v_j'' , and v_k'' . As long as H is below f (i.e., f is inactive), the volume of T_f^{const} does not change when H is moved. The volume, $\text{vol}(T_f^{\text{const}})$, of T_f^{const} can, in fact, be computed using Equation (12) in Lemma 4.1, evaluated at $h = z_i$.

The polyhedron T_f^{var} is a prism with base $\langle v_i, v_k'', v_j'' \rangle$, and height $z_i - h$. The volume, $\text{vol}(T_f^{\text{var}})$, of T_f^{var} can be computed as




Figure 8: The support polyhedron $T_f = T_f^{const} \cup T_f^{var}$ contributed by an inactive back facet f . T_f^{const} is defined by the vertices v_i, v_j, v_k, v'_j , and v''_k . T_f^{var} is defined by $v_i, v'_i, v'_j, v'_k, v''_j$, and v''_k .

$$\text{vol}(T_f^{\text{var}}) = \text{area}(\langle v_i, v_k'', v_j'' \rangle) \cdot (z_i - h),$$

where $\text{area}(\langle v_i, v_k'', v_j'' \rangle)$ is just the area of the vertical projection of f onto the plane $z = z_i$, and depends only on the coordinates of the vertices of f . (It can also be computed using Equation (11) in Lemma 4.1, evaluated at $h = z_i$.)

Let us denote by $\text{inactive-volume}_f(h)$ the support volume contributed by an inactive facet f , w.r.t. the plane $H : z = h$. Then

$$\begin{aligned} \text{inactive-volume}_f(h) &= \text{vol}(T_f^{\text{var}}) + \text{vol}(T_f^{\text{const}}) \\ &= \text{area}(\langle v_i, v_k'', v_j'' \rangle) \cdot (z_i - h) + \text{vol}(T_f^{\text{const}}) \\ &= c_f h + d_f, \end{aligned} \tag{13}$$

for some coefficients c_f and d_f that depend only on the coordinates of the vertices of f .

Finally, if f is a front facet, then a symmetric discussion applies.

The lemma follows by summing the volume contributions of all inactive facets using Equation (13). ■

4.1 The algorithm

The algorithm is similar to the one in Section 3.1. We sweep the plane H upwards, visiting the vertices in sorted order and at each vertex update the active-volume and inactive-volume terms based on the facets incident to the current vertex. The three cases outlined in the algorithm in Section 3.1 hold here as well, except that all references to the active-area and inactive-area terms are now to the active-volume and inactive-volume terms respectively. At each vertex we minimize the sum of the active-volume and inactive-volume terms (hence the total support volume) as the position, h , of H ranges between the z -coordinates of the current vertex and the next one. This gives the optimal position for H between successive vertices. We repeat this for all vertices to find the globally-best position for H .

Theorem 4.1 *The support volume version of Problem 2.1 can be solved in $O(n \log n)$ time for a convex polyhedron \mathcal{P} with n vertices ($O(n)$ time if the vertices are given in sorted order in the decomposition direction \mathbf{d}).*

Proof Similar to the proof of Theorem 3.1. ■

4.2 Experimental results

We have also implemented the above support volume minimization algorithm. The details of the implementation and the test polyhedra are the same as in Section 3.2. Tables 3 and 4 summarize our results.

Support Volume				
#verts n	min volume	h_{min}	run time (s)	non-decomp. volume
20000	2.7	-.17	4.1	1046037.3
40000	1.0	0.02	8.2	1047241.1
60000	0.6	-.08	12.6	1047256.4
80000	0.3	-0.04	16.7	1047234.3
100000	0.3	0	20.9	1047201.5
120000	0.2	.03	25.5	1047197.6
140000	0.2	-.02	29.7	1047167.5
160000	0.1	0	34.1	1047192.4
180000	0.1	0.02	37.9	1047188.5
200000	0.1	-.03	43.2	1047200.0

Table 3: Minimum support volume for convex polyhedra generated from random points on a sphere of radius 100. Here “non-decomp. volume” refers to the support volume when the polyhedra are built without decomposition; observe the significant reductions achieved via decomposition. Note that the support volume without decomposition is roughly equal to the volume under the lower hemisphere, as it should be.

5 Decomposing non-convex polyhedra

We assume hereafter that \mathcal{P} is a non-convex polyhedron. Such polyhedra pose problems that are not encountered in the convex case. First, unlike the convex case, supports need not extend all the way to the platform, but instead may terminate at some other point on the polyhedron itself. Furthermore, it is possible that only part of a facet is in contact with supports, unlike the convex case where either the entire facet or no part of it is in contact with supports. Finally, it is possible that parallel facets are also in contact with supports, either fully or partially. Fortunately, despite these problems, it is possible to handle a non-convex polyhedron in a fashion similar to the convex case, after doing some initial processing on it.

For now, we ignore the issue of controlling the number of pieces in the decomposition; we discuss this in Section 5.7. Furthermore, for concreteness, we focus on the contact-area version of Problem 2.1. (In fact, the volume problem is easier, because parallel facets do not contribute to support volume—as they do to contact-area—and, therefore, can be essentially ignored.)

5.1 Overview of approach

The main idea is to partition each front or back facet of \mathcal{P} into two classes of triangles, called black and gray triangles. (One of these classes may be empty.) A black triangle t will always be completely in contact with supports, regardless of the position of the decomposition plane H ; therefore, it always contributes a fixed amount, $area(t)$, to the total contact-area and so can be ignored for minimization purposes. However, a gray triangle t will contribute anywhere between zero and $area(t)$ to the total contact-area, depending on the position of H , and so needs to be accounted for.

Parallel facets are treated a little differently. A parallel facet of \mathcal{P} is partitioned into three

Support Volume					
#verts n	angle	min volume	h_{min}	mean run time (s)	non-decomp. volume
20001	200	10999.5	-21.1	4.2	965039.4
	37	26389.1	21.5		1107383.9
40001	40	24003.9	21.2	8.6	947591.6
	75	3631.2	-4.4		994452.0
60001	240	12900.6	-4.1	12.7	815518.8
	112	8118.6	2.4		865432.8
80001	80	1359.2	-4.2	17.3	1014386.0
	150	20249.8	-24.6		887435.1
100001	280	1349.3	-4.2	21.5	1014432.7
	187	1068.5	-10.7		1033711.3
120001	120	12892.3	-4.1	25.8	815171.4
	225	13784.8	-19.8		767197.8
140001	320	22932.9	21.4	30.3	948013.1
	263	549.0	3.2		959299.2
160001	160	10760.1	-21.2	34.7	962983.8
	300	12908.9	4.0		849211.7
180001	0	0	1.7	39.5	2128798.5
	338	14356.3	21.2		1754483.3
200001	200	10749.7	-21.2	43.6	962390.9
	15	6736.2	17.2		1952343.8

Table 4: Minimum support volume for convex polyhedra generated from random points on a cone; each polyhedron has been rotated by the indicated angle to make it non-symmetric about the origin. Here “non-decomp. volume” refers to the support volume when the polyhedra are built without decomposition; observe the significant reductions achieved via decomposition.

classes of triangles, called black, gray, and white. (Up to two of these classes may be empty.) Black and gray triangles have the same interpretation as above. A white triangle will never be in contact with supports, regardless of the position of H , and so can be ignored for minimization purposes.

From the above discussion it is clear that only gray facets are relevant to the minimization problem. As we will see, these can be handled using the approach in Section 3.

5.2 Black, gray, and white triangles

We next define formally the different types of triangles and establish their stated properties.

Let f be a front facet. Imagine that we build \mathcal{P} , without decomposition, in direction \mathbf{d} . Consider the supports (if any) that are in contact with f . Their *footprint* on f , i.e., their intersection with f , is a collection of polygons, which we will call *black polygons*. We triangulate the black polygons of f to obtain the set, B_f , of *black triangles* associated with f . Let p be any point in some black triangle. Since p is in contact with supports for build direction \mathbf{d} , the ray r_p , emanating from p in this direction must intersect \mathcal{P} . If p' is the first intersection of r_p and \mathcal{P} —not counting p —then the segment $\overline{p'p}$ is the support for p' . (Informally, p is in contact with supports because some part of \mathcal{P} , i.e., p' , is directly above p in direction \mathbf{d} .) The complement of the black polygons on f is a collection of *gray polygons*, and the set, G_f , of *gray triangles* associated with f is obtained by triangulating the gray polygons. No point in a gray triangle is in contact with supports for build direction \mathbf{d} .

Now, suppose that we build \mathcal{P} with decomposition. Let t be a black triangle in B_f . Assume that the decomposition plane H intersects t , and let t^- and t^+ be the parts of t above and below H , respectively. That is, $t^- \in \mathcal{P}^-$ and $t^+ \in \mathcal{P}^+$. Now, all of t^- will require support since it belongs to the front facet f , and \mathcal{P}^- is built in direction $-\mathbf{d}$. The polyhedron \mathcal{P}^+ is built in direction \mathbf{d} , and, furthermore, the part of \mathcal{P}^+ that is directly above t^+ is the same as the part of \mathcal{P} that would be directly above t^+ if \mathcal{P} were built without decomposition in direction \mathbf{d} . Since t is a black triangle, t^+ is in contact with supports for direction \mathbf{d} in the case without decomposition, and, hence, also in the case with decomposition. Thus, all of t is in contact with supports.

If H is completely above t , then $t^+ = \emptyset$ and $t^- = t$; if H is completely below t , then $t^- = \emptyset$ and $t^+ = t$. In either case, the above argument still applies. It follows that regardless of the position of H , if t is a black triangle, then all of it is in contact with supports. Thus, black triangles may be ignored.

On the other hand, suppose that t is a gray triangle in G_f , where f is a front facet. Assume that H intersects t . As in the case of a black triangle, the part $t^- \in \mathcal{P}^-$ will be in contact with supports. Consider the part $t^+ \in \mathcal{P}^+$, which is built in direction \mathbf{d} . No part of t^+ will be in contact with supports, since in the case where \mathcal{P} is built in direction \mathbf{d} , without decomposition, no part of t is in contact with supports. If H is completely below or completely above t , then either t is not in contact with supports at all or is completely in contact with supports. Therefore, it follows that t 's contribution to the total contact-area depends on the position of H . In fact, this contribution is exactly the same as for a front facet in the convex case. That is, if t is inactive, then its contribution is zero or $area(t)$; if t is active, then its contribution is a quadratic function of the position, h , of H , as in Section 3.

So far, we have considered the case where the facet $f \in \mathcal{P}$ is a front facet. If f is a back facet of \mathcal{P} then we define the black and gray triangles on f by building \mathcal{P} , without decomposition, in direction $-\mathbf{d}$. Then a discussion symmetric to the one above again shows that, in the case with

decomposition, only gray triangles need be considered.

Finally, we turn to the case where f is a parallel facet. Suppose that we build \mathcal{P} without decomposition in direction \mathbf{d} , and, independently, in direction $-\mathbf{d}$. The *black polygons* consist of all points on f that are in contact with supports for both directions. The *gray polygons* consist of all points on f that are in contact with supports for exactly one of the two directions. The *white polygons* consist of all points on f that are not in contact with supports for either of the two directions. By definition, these polygons partition f . The set B_f (resp. G_f, W_f) of *black* (resp. *gray, white*) triangles on f is obtained by triangulating the black (resp. *gray, white*) polygons.

Now suppose that we build \mathcal{P} with decomposition. It is easy to see that regardless of the position of the decomposition plane H , a black triangle will always be in contact with supports, while a white triangle will never be in contact with supports. Therefore, both these types of triangles may be ignored. However, a gray triangle will contribute to the contact-area an amount which depends on the position of H . For instance, let $t \in G_f$ be a gray triangle which exists because f is in contact with supports when \mathcal{P} is built, without decomposition, in direction \mathbf{d} . Then, in the case with decomposition, t^+ is in contact with supports and t^- is not. So the contribution of t to the total contact-area varies quadratically with the position, h , of H , from zero to $area(t)$. So, again only gray triangles need to be considered.

5.3 Computing black and gray triangles for front and back facets

We discuss how to compute the footprint of the supports on each front facet, hence the black and gray triangles on it.

We will use a technique called *cylindrical decomposition* [Mul93]. From each edge, e , of each back facet, b , we erect a strip, $V_{e,b}$, which passes exactly through e and extends vertically downwards, in direction $-\mathbf{d}$. As soon as a part of $V_{e,b}$ intersects another facet of \mathcal{P} (which must be a front facet), we stop propagating that part below the intersected facet; however, we continue propagating the remaining parts of $V_{e,b}$.⁷ Each such intersection of $V_{e,b}$ with a front facet will be a part of the footprints that we are trying to compute.

To perform this step efficiently, we compute the intersection of each front facet with $V_{e,b}$ to get a set, L , of line segments. (Since the facets of \mathcal{P} do not intersect each other, the segments in L are non-crossing, but possibly touching.) We do a trapezoidal decomposition [dBvKOS97] of $L \cup \{e\}$ in the plane of $V_{e,b}$. We identify each trapezoid in this decomposition that is adjacent to e at the top and to a line segment $\ell \in L$ at the bottom. The bottom edge of this trapezoid is one of the sought intersections of $V_{e,b}$ with a front facet. We store this edge with the front facet that generated ℓ .

However, not all footprints on front facets will be discovered by the above process. For instance, if the projection of b completely covers a front facet, f , below it, then none of the strips $V_{e,b}$ erected from b will intersect f , and, yet, supports for b will rest on f . To handle such situations, we also erect from each edge, e , of each front facet, f , a strip $V_{e,f}$ vertically upwards, stopping the propagation of any part of the strip as soon as it intersects a (back) facet above it, while continuing to propagate other parts. To compute these intersections we do a trapezoidal decomposition of the set $L' \cup \{e\}$, where L' contains the intersections of all the back facets with $V_{e,f}$. For each trapezoid in this decomposition that is incident to e at the bottom and to a segment $\ell' \in L'$ at the top, we take the top edge of the trapezoid as the intersection of $V_{e,f}$ with the back facet that generated ℓ' , and store it with that back facet.

⁷The process resembles water cascading down from e .

After we have carried out the above steps for all front and back facets, we have associated with each facet a list of line segments corresponding to intersections of the different vertical strips with the facet. Since a strip is not propagated below an intersected facet, it is easy to see that the line segments associated with a facet are non-crossing (but may be touching). For each facet, we compute the arrangement [dBvKOS97] of the set consisting of the associated segments and the edges bounding the facet.

Let f be any front facet and let c be any cell of the arrangement computed on f . Then c is the footprint of a support on f for build direction \mathbf{d} , hence a black polygon, iff there is a cell c' on a back facet b above f , such that c' projects exactly to c . (The cells c and c' form the bottom and top facets of a support cylinder; the other facets of this cylinder are vertical and bounded below and above by edges of c and c' .) Any other cell of f is a gray polygon.

We can identify the black triangles of f directly (instead of first computing the black polygons). We project the arrangement on each front facet to the xy -plane, triangulate its cells, and then lift the triangles back to the front facet. We make a list, F , of these lifted triangles along with their centroids, and sort F lexicographically on the x -, y -, and z -coordinates of the centroids, taken in that order. We make a similar sorted list B for the back facets. Note that if a cell, c , from a front facet and a cell, c' , from a back facet form the bottom and top facets of a support cylinder, then c and c' have identical projections on the xy -plane and will, therefore, be triangulated identically (if a deterministic triangulation algorithm is used, e.g., [dBvKOS97, Ch. 3]). It should be clear now that a simultaneous scan of the two sorted lists suffices to identify matching pairs of triangles, where one triangle is from B and the other is from F such that the former is above the latter and projects exactly to it. For each matching pair, the triangle from F is a black triangle on some front facet; all unmatched triangles of F are gray triangles on front facets. (This approach for matching triangles, using a lexicographic sort, is due to [CP90].)

In this way, we can compute all the black and gray triangles for each front facet. A symmetric approach yields the black and gray triangles for the back facets.

Lemma 5.1 *The set of black and gray triangles for all front and back facets of an n -vertex polyhedron \mathcal{P} can be computed in $O(n^2 \log n)$ time.*

Proof Correctness is clear from the discussion above. We establish the running time.

The number of line segments generated by intersecting front facets with any strip $V_{e,b}$ is $O(n)$, so their trapezoidal decomposition has size $O(n)$ and can be computed in time $O(n \log n)$ per strip, hence $O(n^2 \log n)$ for all strips.

The total number of segments (intersections) created on all the front facets of \mathcal{P} by the propagation of any $V_{e,b}$ is $O(n)$. This is because the trapezoidal decomposition has $O(n)$ trapezoids and the segments of interest are the bottom edges of a subset of these trapezoids. Since there are $O(n)$ $V_{e,b}$'s, the total number of segments generated by them on all front facets is $O(n^2)$.

Now, consider a front facet, f , and let m_f be the total number of segments on f taken over all strips $V_{e,b}$. By the above discussion, $\sum_f m_f = O(n^2)$. (Note that m_f could be $\Theta(n^2)$ for an individual facet f . This could happen if, for instance, $\Theta(n)$ back facets are above f and overlap in the form of a trellis. However, $\sum_f m_f$ is still $O(n^2)$.) As noted earlier in the discussion of the algorithm, since a strip is not propagated below an intersected facet, the segments on f are non-crossing; thus, their arrangement has complexity $O(m_f)$. It can be computed and triangulated in time $O(m_f \log m_f)$, and yields $O(m_f)$ triangles. Taken over all front facets, there are $O(\sum_f m_f) =$

$O(n^2)$ such triangles, and they can be generated and sorted in time $O(\sum_f m_f \log m_f) = O(n^2 \log n)$. Likewise for the triangles from all the back facets. The resulting lists of $O(n^2)$ triangles each can be scanned in $O(n^2)$ additional time to determine the black and gray triangles. ■

5.4 Computing black, gray, and white triangles for parallel facets

Let f be a parallel facet. Recall that a black (resp. gray, white) triangle on f is one which is in contact with supports for both (resp. exactly one, neither) of the directions \mathbf{d} and $-\mathbf{d}$. Let V_f be the vertical strip which is in the plane of f and exactly contains it. We may assume w.l.o.g. that no vertex of f is in the interior of V_f . (Each bounding line of V_f contains at least one vertex of f . If there is a vertex in the interior of V_f , we draw a vertical line through it and split f into two facets that each satisfy the assumption.) Let vertex u of f lie on one of the bounding lines of V_f and let vertices v and w lie on the other, with v above w in direction \mathbf{d} ; note that the line segments \overline{uv} and \overline{uw} span V_f . Consider the back facets of \mathcal{P} that either pierce V_f above \overline{uv} , or touch V_f above \overline{uv} and are in the same halfspace of V_f as the outer unit-normal, \mathbf{n}_f , of f . (These are the back facets whose supports are potentially in contact with f when \mathcal{P} is built in direction \mathbf{d} .) The intersections of these back facets with V_f is a set, A , of line segments. We do a trapezoidal decomposition of A in the plane of V_f (note that we do not consider the edges of f when doing the decomposition). Let T be the set of trapezoids in this decomposition that are bounded from above by some segment of A and are unbounded below. Let T_f be the trapezoids obtained by intersecting the ones in T with f . (Strictly speaking, some of the elements in T_f may be triangles, which we take to be degenerate trapezoids.)

Figure 9 illustrates the computation of T_f . Figure 9(a) shows the back facets, b_1, b_2, b_3 , and b_4 , that pierce or touch V_f above \overline{uv} : b_1 and b_4 pierce V_f , b_2 touches V_f and is in the same halfspace of V_f as \mathbf{n}_f , and b_3 touches V_f but is in the halfspace of V_f that does not contain \mathbf{n}_f . The supports needed by b_3 will not be in contact with f , so we ignore it. The set A consists of the line segments a_1, a_2 , and a_4 . Figure 9(b) shows the trapezoidal decomposition of A . The trapezoids of T are t_1, t_2 , and t_4 ; their intersections with f , shown shaded, yields the set T_f .

Symmetrically, we consider the front facets of \mathcal{P} that intersect V_f at or below \overline{uv} . These intersections yield a set, A' , of line segments. We do a trapezoidal decomposition of A' , take the set T' of trapezoids in this decomposition that are bounded from below by some segment of A' and are unbounded above, and intersect these with f to get a set T'_f of trapezoids.

The set of black polygons for f is obtained by taking the intersections of every pair of trapezoids, one in T_f and the other in T'_f . Any part of a trapezoid in T_f or T'_f that is not a black polygon is a gray polygon. The complement of the union of the gray and black polygons on f is the set of white polygons. We can compute these three classes of polygons easily by sorting the vertical edges in T_f and T'_f into two lists and then doing a simultaneous scan of the two lists. Since these polygons are actually trapezoids, we can generate the corresponding triangles easily.

Lemma 5.2 *The set of black, gray, and white triangles for all parallel facets of an n -vertex polyhedron \mathcal{P} can be computed in $O(n^2 \log n)$ time.*

Proof The correctness of the method is clear from the above discussion. For the running time, note that the sets A and A' each have size $O(n)$ and can be computed in $O(n)$ time. Doing the trapezoidal decomposition and computing T_f and T'_f takes $O(n \log n)$ time. Finally, sorting and




Figure 9: Computation of the set T_f (shown shaded) for a parallel facet f .

scanning the lists and computing the different polygons takes $O(n \log n)$ time. Therefore each of the $O(n)$ parallel facets can be handled in time $O(n \log n)$, and the lemma follows. ■

5.5 Decomposing to minimize contact-area

At this point, we have for each front, back and parallel facet of \mathcal{P} , a list of the black, gray, and white polygons. As discussed in Section 5.2, only the gray triangles are relevant when decomposing \mathcal{P} to minimize the contact-area of supports. We store the subdivision defined by the union of the set of gray triangles in a doubly-connected edge list and perform a sweep over them as in Section 3 to compute the optimum decomposition plane H . (Even though the algorithm in Section 3 is for convex polyhedra, the sweep does not depend on convexity *per se*, and so it works for our collection of gray triangles as well.)

As seen above, the set of gray triangles can be computed in time $O(n^2 \log n)$. Also, from the proofs of Lemmas 5.1 and 5.2, the number of gray triangles is $O(n^2)$. The time to sort the vertices of the induced subdivision by z coordinates, in preparation for the sweep, is $O(n^2 \log n)$. During the sweep, each vertex can be processed in time proportional to its degree, which implies a time bound of $O(n^2)$ for processing all vertices, since the subdivision has $O(n^2)$ edges.

As noted at the beginning of Section 5, the volume minimization problem is easier, and essentially the same approach works for this also. We will see in Section 5.7 that the size of the decomposition can be controlled in $O(n \log n)$ time. We conclude:

Theorem 5.1 *The contact-area and support volume versions of Problem 2.1 can be solved in $O(n^2 \log n)$ time for a non-convex polyhedron \mathcal{P} with n vertices.*

5.6 Experimental results

Our primary goal was to investigate the extent of support reduction achievable via decomposition of typical models. Therefore, for convenience, we implemented a simpler, but slower, version of the algorithm described above, for support volume minimization. The algorithm differs mainly in how the footprints are computed—instead of using the earlier algorithm for cylindrical decomposition, we use a somewhat less efficient approach, as follows.

Let f be a fixed front facet and let b be any back facet. We project f and b to the xy -plane and compute the intersection of their projections (i.e., triangles), which yields a convex polygon, $C(b)$. If $C(b) \neq \emptyset$, then let p be any point in it, say the centroid. If the pre-images, p_f and p_b , of p on f and b , respectively, are such that p_b is above p_f in direction \mathbf{d} , then p_f is in contact with supports. This implies that the pre-image $C_f(b)$ of $C(b)$ on f is in contact with supports. This follows since no facet of \mathcal{P} pierces another, so there cannot be another point p' in $C(b)$, whose pre-images on f and b are in the opposite order from those of p . (Note that it need not be the case that the cylinder bounded by $C_f(b)$ and by $C_b(b)$ —the pre-image of $C(b)$ on b —is a support cylinder, since b , or parts thereof, need not be immediately above f ; there could be parts of other back facets in between.) Given the polygons $C(b)$ that are found to be in contact with supports, we can compute the footprint of supports on f by taking the union of the pre-images, $C_f(b)$, of these polygons. (In our implementation, we used the functions provided by LEDA [MN99] to perform the union and intersection operations.) We triangulate the footprint (resp. the complement of the footprint) on f to get the black (resp. gray) triangles on f . We handle back facets in a symmetric fashion, using




Figure 10: Non-convex models tested. Here “z-span” denotes the span of coordinates in the z -direction; it is given to help visualize the position, h_{min} , of the decomposition plane. The decomposition direction, \mathbf{d} , points vertically upwards.

direction $-\mathbf{d}$ instead of \mathbf{d} . Thereafter, we apply the sweep-based algorithm on the set of all gray triangles to compute the optimum position of the decomposition plane.

The most expensive part of this algorithm turns out to be the union step in the computation of the footprints. Note that the algorithm simply projects all back facets down to the xy -plane, without regard to any intervening facets. Thus, the complexity of the union of the polygons $C(b)$ on a single front facet, f , can be $\Theta(n^2)$ in the worst case, and $\Theta(n^3)$ over all front facets. (An example of this is a configuration of $\Theta(n)$ front facets stacked on top of each other and $\Theta(n)$ back facets above them that overlap in the form of a trellis. Notice that this situation would not create a problem for cylindrical decomposition, since there we do not project through intervening front facets.) The total time to compute the union is $O(n^2 \log n)$ in the worst case for any front facet [dBvKOS97], hence $O(n^3 \log n)$ over all front facets.

We performed our experiments on the six models shown in Figure 10. The first five models were obtained from Stratasys, Inc., a Minnesota-based company specializing in LM. To keep the running

	model	# facets	min volume	h_{min}	non-decomp. volume	run time (s)	bounding box $l \times w \times h$
1	<code>0-2190.stl</code>	3,492	0.1	0.0	0.7	16,661	$1.4 \times 1.5 \times 1.4$
2	<code>speedo.stl</code>	2,500	0.9	1.9	5.1	15,730	$2.0 \times 1.9 \times 3.2$
3	<code>mj.stl</code>	2,832	1.7	2.1	8.1	13,911	$2.8 \times 4.6 \times 3.2$
4	<code>r59043b.stl</code>	3,386	0.8	0.4	3.0	20,833	$2.4 \times 2.1 \times 1.4$
5	<code>engine2.stl</code>	4,180	174.9	1.3	251.6	41,156	$12.1 \times 6.2 \times 9.7$
6	<code>cc.stl</code>	112	823,210.0	1.0	823,210.0	4.7	$102.0 \times 102.0 \times 100.0$

Table 5: Experimental results for non-convex models. Experiments were performed on a SUN Ultra 60 Sparc machine with 512 MB of main memory and a 450 MHz processor.

time of our algorithm on these models reasonable, we reduced the number of facets in some of the larger models—specifically, `0-2190.stl` (to about 25% of its original number), `speedo.stl` (to about 15%), and `engine2.stl` (to about 18%). For this we used the *Decimator* software package from Raindrop Geomagic, Inc., a North Carolina-based company specializing in the design of CAD software. This package reduces model sizes while preserving the original topology to the maximum extent possible. The sixth model, `cc.stl`, was hand-generated; it is used to show that on some models our algorithm may not achieve any reduction in support requirements.

Table 5 summarizes our experimental results for the six models. To give an idea of the relative sizes of the models, the table also gives the dimensions of their axes-parallel bounding boxes. The models are listed in the table in decreasing order of the reduction in support volume that is achieved (the sixth column divided by the fourth). On the first four models, our algorithm achieved a reduction ranging from a factor of seven to about a factor of four. On the fifth model, it achieved a reduction of about 1.4. It achieved no reduction at all on the last model, `cc.stl`, which is essentially a hollow cube, with walls of thickness 1 and with four holes on the top face. If this model is built without decomposition, the interior of the cube will be filled with supports, except for the regions directly below the holes. Any decomposition plane that intersects the model strictly above the inner base (which is at height 1), will generate two pieces for which the total support volume is higher than without decomposition, since the region in the lower piece that was originally below the holes is now filled with supports. The optimal position for decomposition is any height in the interval $[0, 1]$.

5.7 Controlling the size of the decomposition

The optimal plane computed by the algorithm in Section 5.5 could decompose a non-convex polyhedron \mathcal{P} into many polyhedra (as many as $\Theta(n)$ in the worst case), which is undesirable since it increases the cost of re-assembling \mathcal{P} . Ideally, the designer should be able to specify an integer K , and the algorithm should compute, among all possible planes that generate no more than K polyhedra, a plane which is optimal w.r.t. support contact-area or volume. We show how this can be done by incorporating a preprocessing step in the algorithm. The idea is to partition the z -axis into $O(n)$ intervals, I_j , such that all planes whose heights are in I_j decompose \mathcal{P} into the same number, k_j , of polyhedra. We then run the sweep algorithm of the previous section but do the minimization step only in those intervals I_j for which $k_j \leq K$.

Let $z_1 < z_2 < \dots < z_t$, $t \leq n$, be the distinct z -coordinates of the n vertices of \mathcal{P} . The

preprocessing involves two sweeps. The first sweep is in the positive z direction and it computes a set of intervals on the z -axis and associates with each interval an integer which is the number of connected components of \mathcal{P}^- generated by any plane whose height is in the interval. Observe that the number of connected components of \mathcal{P}^- w.r.t. a plane of height z_j is the same as the number w.r.t. a plane whose height is anywhere in the interval (z_j, z_{j+1}) ; let this number be k_j^- . Thus, the first sweep computes intervals of the form $[z_j, z_{j+1})$ and associates with each the integer k_j^- . Symmetrically, the second sweep is in the negative z direction and it computes intervals of the form $(z_j, z_{j+1}]$ and associates with each an integer k_j^+ , which is the number of connected components of \mathcal{P}^+ w.r.t. any plane whose height is in $(z_j, z_{j+1}]$. Once these two sets of intervals have been computed, a single scan of them suffices to compute the desired intervals I_j and the corresponding integers k_j . Specifically, each interval I_j is either of the form $[z_j, z_j]$, with $k_j = k_j^- + k_{j-1}^+$, or of the form (z_j, z_{j+1}) , with $k_j = k_j^- + k_j^+$.

We describe the first sweep in more detail. At any time, the vertices of the different connected components of \mathcal{P}^- form a collection of disjoint sets. We maintain these using a Union-Find-Makeset data structure [CLR90]. We initialize the structure to empty and set the current number, c , of connected components of \mathcal{P}^- to zero. Let z_j be the current z -coordinate in the sweep and let V_j be the set of vertices of \mathcal{P} at this z -coordinate. We consider each vertex $v \in V_j$ in turn and process it as follows: We create a new set containing just v and increment c by one. Then for each neighbor, u , of v such that u is already in the Union-Find-Makeset data structure we do the following: If u and v are in different connected components, then we union the sets containing u and v , and decrement c by one. Notice that in this sweep, we only “add” edges to \mathcal{P}^- , so the connected components of \mathcal{P}^- always merge, never split. Thus, a Union-Find-Makeset structure suffices to maintain the connected components of \mathcal{P}^- . After all vertices of V_j have been processed, we set k_j^- to c . At the end of the sweep, all the intervals $[z_j, z_{j+1})$, and their associated integers k_j^- will have been computed.

Excluding the $O(n \log n)$ time for the sorting, the algorithm takes $O(n\alpha(n))$ time, where $\alpha(n)$ is the slow-growing inverse Ackerman function. Hence this preprocessing step does not affect the asymptotic running time of the decomposition algorithm of Section 5.5.

6 Conclusions and future work

In this paper, we have presented a new approach to LM, in which the model is decomposed by a plane into several pieces that can be built independently and then glued together. Several advantages of this approach have been identified, including improved speed, greater versatility, and reduced support requirements. We have presented efficient geometric algorithms to decompose both convex and non-convex models so as to minimize the contact-area and the volume of the supports (Problem 2.1). For convex models, the algorithms are based on the plane sweep paradigm, where formulas for support contact-area and volume are generated, updated, and optimized as a function of the height of the sweep plane. Non-convex models are handled by first identifying certain critical facets that are relevant to the minimization and then using the same approach as in the convex case. Experimental results have been given that demonstrate the utility of the decomposition-based method.

There are actually four versions of Problem 2.1 depending on how we choose to build \mathcal{P}^+ and \mathcal{P}^- : (1) \mathcal{P}^+ and \mathcal{P}^- in directions \mathbf{d} and $-\mathbf{d}$, respectively—the version solved in the paper; (2) \mathcal{P}^+

and \mathcal{P}^- in directions $-\mathbf{d}$ and \mathbf{d} , respectively; (3) \mathcal{P}^+ and \mathcal{P}^- in directions \mathbf{d} and \mathbf{d} , respectively; and (4) \mathcal{P}^+ and \mathcal{P}^- in directions $-\mathbf{d}$ and $-\mathbf{d}$, respectively. For a given model, \mathcal{P} , the support requirements of one of the methods might be better than the other three. For instance, let \mathcal{P} consist of two pyramids, with a common vertex v (i.e., one pyramid is a reflection of the other through v), and let \mathbf{d} be perpendicular to the base of one pyramid. If we choose the decomposition plane H such that it contains v , then method 2 would require no supports. For any choice of H , the other three methods would require a non-zero amount of supports. Given a model, one strategy would be to run all four methods independently and pick the one which minimizes the support requirements.

For convex \mathcal{P} , the solutions for methods 2–4 are similar to that for method 1. For contact-area minimization, method 2 is completely symmetric to method 1. In method 3 (resp. method 4), only back (resp. front) facets need supports, regardless of the position of H ; thus, the support contact-area equals the total area of the back (resp. front) facets, so any position of H is optimal. The situation is similar for support volume minimization, except that the platform for \mathcal{P}^+ and/or \mathcal{P}^- need no longer be H . For instance, in method 2, the platform for \mathcal{P}^+ (resp. \mathcal{P}^-) is the plane parallel to H and coinciding with the extreme vertex in direction \mathbf{d} (resp. $-\mathbf{d}$), so the support volume needs to be computed w.r.t. this platform rather than w.r.t. H . Similar statements apply to methods 3 and 4. In these methods, even though only back or front facets need supports, the position of H does affect the support volume (unlike contact-area); the optimal position for H can be found by minimizing the expression for the total support volume, as in method 1.

Unfortunately, the approach in Section 5 (where facets are decomposed into black, gray, and white triangles) does not extend to methods 2–4 when \mathcal{P} is non-convex. Consider, for instance, method 2. Let t be a black triangle from a front facet and assume that H intersects t . Now, t^+ is completely in contact with supports, since it is built in direction $-\mathbf{d}$. However, the structure of the supports that will be in contact with t^- depends on the (yet-to-be-built) portion of \mathcal{P}^- that is above H , which can be quite complex. (This problem did not arise in method 1, for reasons discussed at the beginning of Section 5.2.) Similar problems arise in method 3 for a black triangle belonging to a front facet and in method 4 for a black triangle belonging to a back facet. At present, we do not know how to overcome these difficulties and it appears that a different approach will be needed.

Throughout the paper, we have assumed a fixed decomposition direction \mathbf{d} and found an optimum plane that is normal to \mathbf{d} . A challenging next step is to consider the problem of computing, over all directions \mathbf{d} , an optimum decomposition plane (while also limiting the number of pieces). The methods developed in this paper could be useful in computing such a plane once a small set of candidate directions has been identified.

Acknowledgement

We thank Stratasys, Inc. for providing us with STL files for our experiments. We also thank the referees for helpful comments that improved the paper substantially.

References

- [ABB⁺97] B. Asberg, G. Blanco, P. Bose, J. Garcia-Lopez, M. Overmars, G. Toussaint, G. Wilfong, and B. Zhu. Feasibility of design in stereolithography. *Algorithmica*, 19:61–83, 1997.
- [AD95] S. Allen and D. Dutta. Determination and evaluation of support structures in layered manufacturing. *Journal of Design and Manufacturing*, 5:153–162, 1995.
- [AD00] P. Agarwal and P. Desikan. Approximation algorithms for layered manufacturing. In *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 528–537, 2000.
- [Bai95] M.J. Bailey. Tele-manufacturing: Rapid prototyping on the Internet. *IEEE Computer Graphics and Applications*, 15(6):20–26, 1995.
- [Bau75] B. G. Baumgart. A polyhedron representation for computer vision. In *Proceedings of the AFIPS National Computer Conference*, volume 44, pages 589–596, 1975.
- [BB95] M. Bablani and A. Bagchi. Quantification of errors in rapid prototyping processes and determination of preferred orientation of parts. In *Transactions of the 23rd North American Manufacturing Research Conference*, 1995.
- [BDH96] C. Bradford Barber, David P. Dobkin, and Hannu Huhdanpaa. The quickhull algorithm for convex hulls. *ACM Transactions on Mathematical Software*, 22(4):469–483, December 1996. (See also <http://www.geom.umn.edu/software/qhull/>).
- [BK98] G. Barequet and Y. Kaplan. A data front-end for layered manufacturing. *Computer-Aided Design*, 30(4):231–243, 1998.
- [Bøh95] J.H. Bøhn. Removing zero-volume parts from CAD models for layered manufacturing. *IEEE Computer Graphics and Applications*, 15(6):27–34, 1995.
- [Bos95] P. Bose. *Geometric and computational aspects of manufacturing processes*. PhD thesis, School of Computer Science, McGill University, Montréal, Canada, 1995.
- [CLR90] T.H. Cormen, C.E. Leiserson, and R.L. Rivest. *Introduction to Algorithms*. MIT Press and McGraw-Hill, 1990.
- [CP90] B. Chazelle and L. Palios. Triangulating a nonconvex polytope. *Discrete & Computational Geometry*, 5:505–526, 1990.
- [dBvKOS97] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf. *Computational Geometry: Algorithms and Applications*. Springer-Verlag, Berlin, 1997.
- [DKPS98] D. Dutta, V. Kumar, M. Pratt, and R. Sriram. Towards STEP-based data transfer in Layered Manufacturing. In *Proceedings of the Tenth International IFIP WG 5.2/5.3 Conference PROLOMAT*, 1998. (See also <http://www.mel.nist.gov/msidlibrary/summary/9826.html>).

- [DM94] A. Dolenc and I. Mäkelä. Slicing procedures for layered manufacturing techniques. *Computer-Aided Design*, 26:119–126, 1994.
- [FF94] D. Frank and G. Fadel. Preferred direction of build for rapid prototyping processes. In *Proceedings of the 5th International Conference on Rapid Prototyping*, pages 191–200, 1994.
- [FM97] S. Fekete and J. Mitchell. Histogram decomposition and stereolithography, 1997. Manuscript. <http://www.zpr.uni-koeln.de/~paper/paper.php3?paper=280>.
- [Gol91] R.N. Goldman. Area of planar polygons and volume of polyhedra. In J. Arvo, editor, *Graphics Gems II*, pages 170–171. Academic Press, 1991.
- [Jac92] P. Jacobs. *Rapid Prototyping & Manufacturing: Fundamentals of StereoLithography*. McGraw-Hill, 1992.
- [Joh99] E. Johnson. Support generation for three-dimensional layered manufacturing. Master’s project report, Dept. of CS&E, Univ. of Minnesota, Minneapolis, MN, 1999.
- [KD96] P. Kulkarni and D. Dutta. An accurate slicing procedure for layered manufacturing. *Computer-Aided Design*, 28:683–697, 1996.
- [Maj98] J. Majhi. *Geometric methods in computer-aided design and manufacturing*. PhD thesis, Dept. of Computer Science & Engineering University of Minnesota, Minneapolis, MN, 1998.
- [MJS⁺99] J. Majhi, R. Janardan, J. Schwerdt, M. Smid, and P. Gupta. Minimizing support structures and trapped area in two-dimensional layered manufacturing. *Computational Geometry: Theory & Applications*, 12:241–267, 1999.
- [MJSG99] J. Majhi, R. Janardan, M. Smid, and P. Gupta. On some geometric optimization problems in layered manufacturing. *Computational Geometry: Theory & Applications*, 12:219–239, 1999.
- [MJSS98] J. Majhi, R. Janardan, J. Schwerdt, and M. Smid. Multi-criteria optimization algorithms for layered manufacturing. In *Proceedings of the 14th Annual ACM Symposium on Computational Geometry*, pages 19–28, 1998.
- [MKDP98] A. Marsan, V. Kumar, D. Dutta, and M. Pratt. An assessment of data requirements and data transfer formats for Layered Manufacturing. Technical Report NISTIR 6216, National Institute of Standards and Technology, September 1998. (See also <http://www.mel.nist.gov/msidlibrary/summary/9821.html>).
- [MN99] K. Mehlhorn and S. Näher. *LEDA: A Platform for Combinatorial and Geometric Computing*. Cambridge University Press, Cambridge, U.K., 1999.
- [MS99] S. McMains and C. Séquin. A coherent sweep plane slicer for layered manufacturing. In *Proceedings of the Fifth ACM Symposium on Solid Modeling and Applications*, pages 285–295, 1999.

- [Mul93] K. Mulmuley. *Computational Geometry: An introduction through randomized algorithms*. Prentice–Hall, 1993.
- [SBE95] P. Stucki, J. Bresenham, and R. Earnshaw. Computer graphics in rapid prototyping technology. *IEEE Computer Graphics and Applications*, 15(6):17–19, 1995. Guest Editors’ Introduction.
- [SSJ+00] J. Schwerdt, M. Smid, R. Janardan, E. Johnson, and J. Majhi. Protecting critical facets in layered manufacturing. *Computational Geometry: Theory & Applications*, 16:187–210, 2000.
- [SSMJ99] J. Schwerdt, M. Smid, J. Majhi, and R. Janardan. Computing the width of a three-dimensional point-set: an experimental study. *ACM Journal of Experimental Algorithmics*, Vol. 4, Art. 8, 1999. Available online at <http://www.jea.acm.org/1999/SchwerdtWidth/> .