2.3

Removing Small Features with Real CAD Operations

Brett W. Clark

Sandia National Laboratories
P.O. Box 5800 MS 0376
Albuquerque, New Mexico
87185-0376
bwclark@sandia.gov

Abstract

Preparing Computer Aided Design models for successful mesh generation continues to be a crucial part of the design to analysis process. A common problem in CAD models is features that are very small compared to the desired mesh size. Small features exist for a variety of reasons and can require an excessive amount of elements or inhibit mesh generation all together. Many of the tools for removing small features modify only the topology of the model (often in a secondary topological representation of the model) leaving the underlying geometry as is. The availability of tools that actually modify the topology and underlying geometry in the boundary representation (B-rep) model is much more limited regardless of the inherent advantages of this approach. This paper presents a process for removing small featrues from a B-rep model using almost solely functionality provided by the underlying solid modeling kernel. The process cuts out the old topology and reconstructs new topology and geometry to close the volume. The process is quite general and can be applied to complex configurations of unwanted topology.

1 Introduction

Preparing CAD models for mesh generation involves various processes which may include CAD format translation, geometry and topology generation and repair, defeaturing, and decomposition into meshable pieces. Defeaturing is usually required because the meshing algorithms rely heavily on the topology and underlying geometry of the boundary representation (B-rep) model and in many cases the B-rep model resulting from design does not meet the stringent requirements of the meshing algorithm. Of particular interest in this work is the removal of features that are very small compared to the desired mesh size.

These features require unnecessary resolution in the mesh or inhibit mesh generation completely.

Because B-rep models used for meshing can originate from a number of different sources, small features also have many different origins. Small features can be actual features that were intended in the design but which have little meaning in the analysis. They can also be features that were not intended in the design but which were a result of a careless CAD designer. During design, sliver surfaces and curves are often introduced by the solid modeling kernel to ensure a water-tight volume. Furthermore, small features can result from translating a model from a CAD format with loose tolerances to a CAD format with tight tolerances. As the analyst usually does not have control over these issues it is crucial to have tools for eliminating small features.

One place to eliminate small features is in the native CAD system where design took place. There, unwanted CAD features can simply be suppressed or modified directly. However, CAD systems do not generally provide many tools for fixing small topology that is not one of the provided CAD features. Therefore, the analyst is forced to use whatever geometry cleanup functionality is available in downstream applications. As a result, any modifications to the model will generally be done to the topology at the B-rep level and not at the feature level. Most defeaturing capabilities in applications downstream of the CAD design system fall into 1 of 3 categories: 1) modify the topology and geometry of the B-rep model directly using solid modeling or "real" operations if available, 2) modify a virtual representation of the topology in the B-rep model using "virtual" topology operations [6], and 3) modify the generated mesh to remove adverse affects of unwanted features in the B-rep model.

This work presents a small feature removal process that removes unwanted topology and associated geometry from a B-rep model. The proposed process works directly on the B-rep model reconstructing the geometry and topology in the vicinity of the feature(s) being removed using a combination of real operations provided by the underlying solid modeling kernel and virtual operations provided by the meshing application. The implementation was done in the CUBIT mesh generation package [1] developed at Sandia National Labs which supports the ACIS solid modeling kernel [2]. The solid modeling kernel operations employed and referred to during the description of the algorithm are operations provided by ACIS. However, the process could be implemented with any solid modeling kernel.

2 Related Work

As stated in the introduction, most defeaturing capabilities in applications downstream of the CAD design system fall into 1 of 3 categories: 1) modify the topology and geometry of the B-rep model directly using solid modeling or "real" operations if available, 2) modify a virtual representation of the topology in the B-rep model using "virtual" topology operations (composites/merges and partitions/splits), and 3) modify the generated mesh to remove adverse affects of unwanted features in the B-rep model.

The first approach, modifying the B-rep model directly with real operations, has the advantage that modifications made can be saved in the native solid modeling kernel format for later use in other applications. These types of modifications also have the advantage that in most cases not only is the topology modified but the underlying geometric definitions are modified maintaining a consistent relationship between the topology and geometry. Despite these advantages, however, the breadth of tools that modify the B-rep topology and geometry directly is limited. Research in this area includes work by Eccles and Steinbrenner [3]. They describe tools for generating B-rep topology information for models that don't provide it (some IGES files for example). The tools include algorithms for merging curves within a given tolerance and reconstructing single surfaces from networks of surfaces and curves. The main benefit of these tools appears to be in generating topology information but could probably also be extended to do general defeaturing tasks. These capabilities all appear to be for surface meshing applications. Butlin and Stops [4] describe tools for fixing common problems resulting from CAD format translation. They also mention defeaturing capabilities such as "joining patchworks of faces and chains of edges" and "collapsing of edges and faces" but do not go into much detail about these capabilities. Similarly, Jones et al. [5] describe tools for general gap cleanup and defeaturing including tools to make a single NURBS curve out of a chain of curves. Analogous capabilities for surfaces are discussed but not yet provided. Lee et al. [6] describe a system for representing feature modification. The system is independent of whether the implementation is done with real B-rep operations or virtual operations but they give some examples using the former. This work is focused more on the theory of a system that can provide reversible operations and not on the introduction of new operators. Venkataraman and Sohoni [7] present a delete face operator that is similar to the delete face operators provided by commercial solid modeling kernels but which handles a wider variety of cases. This operator appears to be quite useful for cases where the faces neighboring the deleted faces can intersect to close the volume.

The next area of research is very similar to the first except the topology operations are applied to a secondary representation of the B-rep topology. This secondary representation is sometimes referred to as virtual topology and allows for topological modifications without modifying the underlying geometry. Various incarnations of this approach have been implemented [6, 8, 9, 10, 11, 12]. The topological operations provided are generally based on two main operators: composite or merge and partition or split. With these two building blocks various topological modifications can be accomplished. As powerful as these tools have become there are still some limitations. First, the topology modifications can usually only be used in the application that defines them. In most cases they can not be saved out with the B-rep model to be used in other applications. Second, just modifying the topology and not the underlying geometry can lead to difficulties in geometric evaluations (surface normals, for example) in downstream algorithms such as mesh generation. Third, virtual topology operators are generally not interoperable with real solid modeling operations. This means that real solid modeling operations usually cannot be applied after virtual operations have been applied. The reason for this is that the virtual operations are typically applied to a secondary layer of topology that sits on top of the original B-rep model. As a result, the solid modeling kernel knows nothing about the virtual topology and real operations often invalidate the virtual topology. At the time of this writing, however, the author is aware of yet unpublished advancements in interoperability between real and virtual operators made in the CUBIT mesh generation toolkit developed at Sandia National Labs.

The third area of research is somewhat different from the first two in that modifications are made to the mesh after mesh generation and not to the original B-rep model. Eccles and Steinbrenner [3] use this approach to help generate model topology from meshed surfaces. The meshed surfaces are "stitched" together after modifying the mesh at the junction. Mobley et al. [12] first perform curve meshing to help guide them through the defeaturing process. The surface meshing then follows using the modified curve meshes. Dey et al. [13] first generate a mesh and then use mesh quality to determine where to modify the mesh to remove the adverse effects from small features in the geometry. In a follow-on work [14] Shephard et al. took this approach a step further and modified the topology (with virtual operations) based on modifications to the mesh.

3 Small Feature Removal Process

The small feature removal process presented here has two main parts: 1) small feature recognition and preprocessing via a surface splitting algorithm and 2) small feature removal via the "remove topology" operator. Both of these will now be described.

3.1 Small Feature Recognition and Preprocessing

Prior to removing unwanted small features in the B-rep model it is often necessary to examine the neighboring topology and make modifications that will facilitate the removal process. This is especially true when removing complex configurations of small features. This section describes a surface splitting algorithm that was developed for this purpose.

Small features can show up in a number of different ways. The simplest types to detect are small curves and small surfaces where the curve length and surface area can be calculated and compared to a reference value. However, small features can also exist when topological entities come in close proximity to one another. For example, the surface in Fig. 1 necks down to a very narrow region even though there is not a small curve length or small surface area to detect. The surface splitting algorithm developed as a part of this research attempts to find narrow regions of surfaces that can be split off into individual surfaces. These splits are "real" splits done using the solid modeling kernel. Fig. 1 shows how the algorithm would split the current surface into three

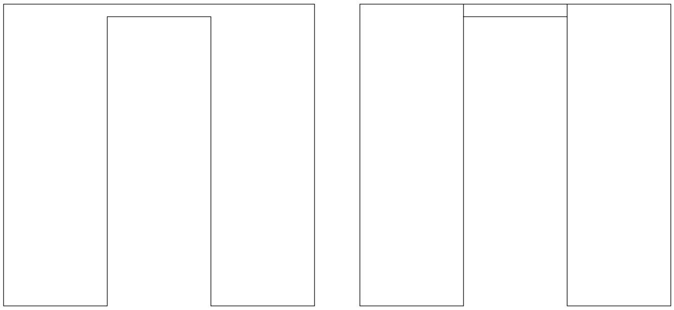

Fig. 1. Surface with narrow region

different surfaces. The algorithm checks the distance from endpoints of curves to other curves to find close proximities. When endpoints are found to be close to other curves, additional points along the curves are compared to locate the extents of the close regions. There are cases that this approach will not catch (like curves coming into close proximity to one another at their mid sections) but it has proven sufficient for the current needs. Fig. 2 shows an example of a model that has been modified by the surface splitting algorithm. The effect is to break up surfaces with narrow regions into the simplest set of surfaces based on surrounding topology. This usually results in reducing the number of curves in a surface to 3 or 4.

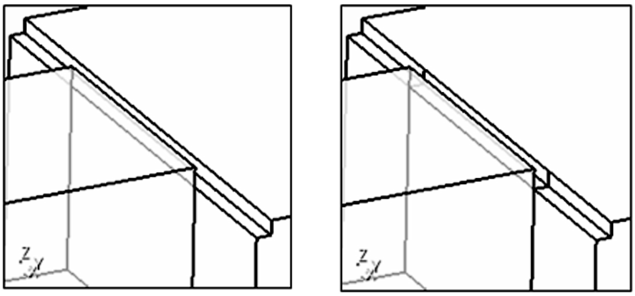

Fig. 2. Before and after performing automatic splitting of surfaces

3.2 Small Feature Removal

After the surface splitting algorithm has been applied the “remove topology” operator is used to remove the unwanted small features from the B-rep model. This is done using a combination of functionality provided by the underlying solid modeling kernel and virtual composites. As all of the critical modifications to topology and geometry are done directly to the B-rep model using the solid modeling kernel, the above-mentioned advantages of “real” operations are realized with this operator.

The remove topology operation has four main stages: 1) construct new topology and geometry to replace the “old topology” (topology being removed), 2) locally “cut” out the old topology, 3) “stitch” in the new topology, and 4) clean up extraneous curves. Each of these stages will be discussed below.

New Topology/Geometry Construction

The remove topology operator is similar to the remove surface operator provided in many solid modeling kernels. A remove surface operator generally removes a surface from the model and then reintersects adjacent surfaces to create a water-tight topology. However, as shown in Fig. 3, this approach will not always work. In the model shown in Fig. 3 there is a small step represented with a surface. The narrow surface may be removed from the model, but the adjacent surfaces will not intersect to close the model because they are parallel.

The remove topology operator is different from the remove surface operator in that it does not rely on reintersecting the adjacent surfaces to close the topology. The first part of the remove topology operation is to determine what “new topology” to use to replace the “old topology” (topology being removed). In the example in Fig. 3 the remove topology operator will generate new topology to replace the narrow surface representing the step. The remove

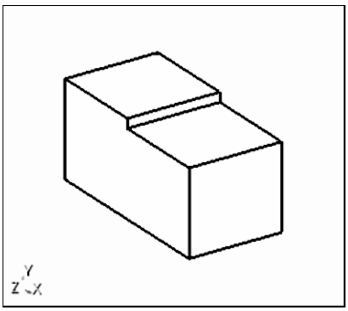

Fig. 3. Model with a small step surface

topology operator tries to reduce the dimension of the old topology wherever possible. In this example it will try to reduce the narrow step surface to a single curve as depicted by dotted lines in Fig. 4. The algorithm for doing

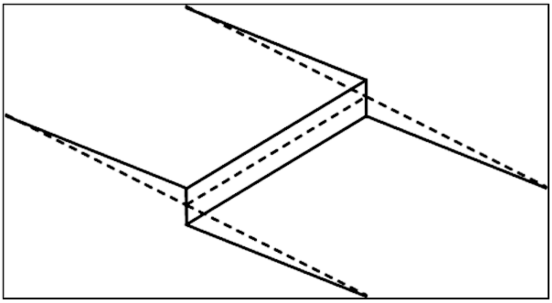

Fig. 4. Dimension reduction of narrow surface

this first looks at all of the topology specified for removal and finds vertices that are closer than some characteristic small curve length. These vertices are "clumped" together. Each clump of vertices will result in a single new vertex in the new topology. The clumps for the example above are shown in Fig. 5. The second step is to decide what new topology will exist between the

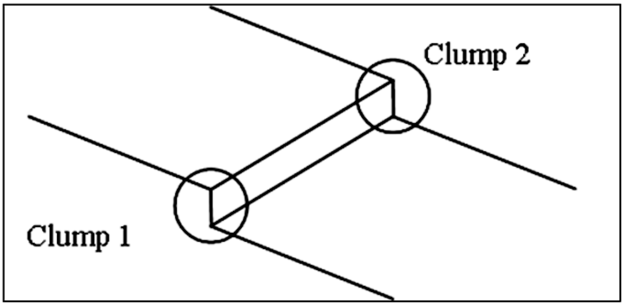

Fig. 5. Clumps of vertices in original topology

clumps. The old topology is examined to see how vertices were connected with curves. Curves that were between vertices in the same clump will be removed from the model as they are smaller than the characteristic small curve size. As a rule, in the new topology, clumps will be connected by a single curve. Therefore, surfaces whose vertices are completely contained in two clumps, as in the example in Fig. 5, will be removed and replaced by the new curve connecting the two clumps. Multiple surfaces can be removed simultaneously in this manner as shown in Fig. 6. Because the algorithm will handle networks of small features at the same time it is often advantageous to use the automatic

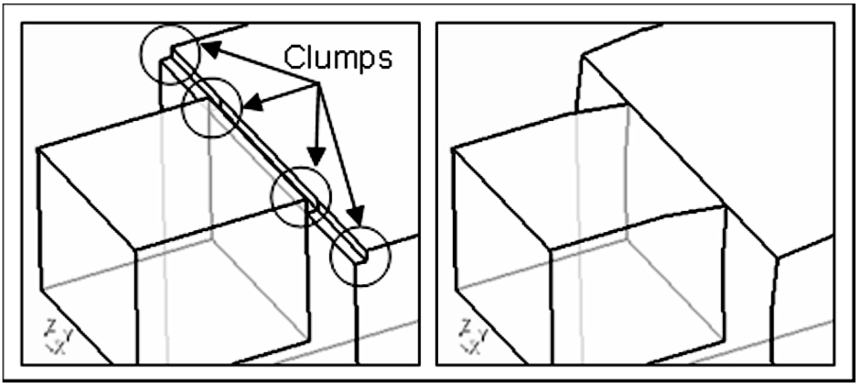

Fig. 6. Multiple surfaces removed between clumps

surface splitting algorithm described above prior to doing the remove topology operation. For example, the model on the left in Fig. 6 had some of its surfaces split before applying the remove topology operator. Fig. 2 shows this model before and after the splits. This splitting process produces vertices and curves that contribute to and simplify the remove topology operation.

Part of the difficulty in generating the new topology is deciding where to put the vertex representing a clump and how to define the geometry of the curves that connect clumps. There are various factors that could go into these decisions. Currently, the position of the new clump vertex is chosen as the average position of all of the vertices in the clump. As there are often multiple curves connecting clumps as in the example in Fig. 6, there can be multiple choices for the new curve connecting the clumps. Similar to the averaging approach for the vertices the new curve between two clumps is a spline defined by average positions along the old curves connecting the clumps. This choice for defining the geometry underlying the new topology is arbitrary and may not be the best choice in some situations. One alternative that would improve the flexibility of the process would be to give the user options for controlling the creation of the geometry.

"Cutting" out the Old Topology

The second part of the operation is to locally cut out the old topology. "Locally" is defined by a user-defined distance representing the distance that will be backed off from the old topology before making the "cuts." Each curve connected to the old topology is split at the specified distance and then the surfaces attached to these curves are split so that the entire network of old topology can be removed from the boundary representation of the volume. These split operations are done using common curve and surface splitting functionality provided by the solid modeling kernel. Fig. 7 shows an exploded view of the old topology being cut out of the boundary representation of the model. The cuts are made in such a way as to facilitate the use of general four-sided surfaces when stitching the new topology into the model. As a result the algorithm must examine the topology and geometry around the old topology and determine an optimal number of four-sided surfaces to use. The cuts are then made based on this analysis.

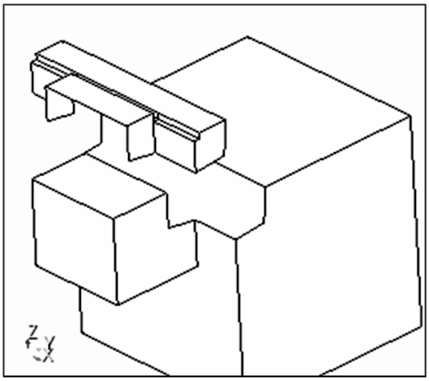

Fig. 7. Exploded view of old topology being cut out of boundary representation

"Stitching" in the New Topology

The last part of the operation is to "stitch" the new topology into the model. In the first part of the operation the new topology and corresponding geometry were defined. However, curves and surfaces to interface this new topology with the rest of the model were not defined. Four-sided (and on rare occasions three-sided) surfaces are used to connect the new topology to the rest of the model. In some cases the four curves bounding the surface will lead to a planar surface, but for all other cases four-sided splines are used. Most solid modeling kernels provide functionality for generating a spline surface from three or four bounding curves. Four-sided surfaces were chosen for their robustness when doing the final stitching to produce a water-tight volume. However, there is no limitation in the algorithm that requires four-sided surfaces. In fact any combination of surfaces could be used to connect the new topology to the rest of the model as long as the result is a water-tight volume. Fig. 8 shows the new topology stitched into the volume. The orange curves (new topology) were generated during the "new topology generation" stage of the operation and the red faces are the four-sided splines used to interface the new topology with the rest of the model. The solid modeling kernel's stitch operator was used to stitch everything together into a water-tight volume. A conscious decision was made to make the modifications to the model as local to the topology being removed as possible. One reason for this is to minimize the amount

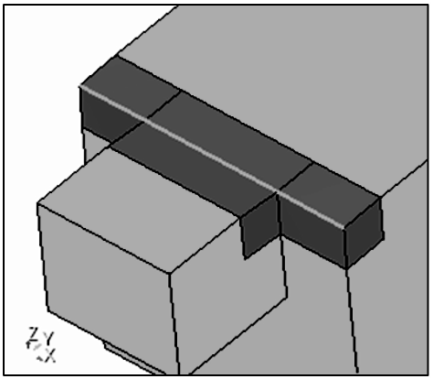

Fig. 8. New topology stitched to volume

of the model approximated by the four-sided surface patches. The interior of these surfaces is completely defined by their bounding curves. Therefore, these surfaces are not good for replacing other surfaces with lots of curvature unless the curvature is represented in the bounding curves.

Cleaning up Extraneous Curves

At this point in the operation the volume is water tight and the only thing left to do is get rid of the curves that were introduced by cutting out the old topology and stitching in the new topology. This step is optional and may not be critical depending on the downstream application. For example, if the downstream application is tetrahedral meshing the curves may not need to be removed depending on the mesh size requirements. The user has some control over how large a region is cut out when removing topology so that if the cutout region size is close to the mesh size there probably won't be a need to remove the curves. If the downstream application is hexahedral meshing it may be necessary to remove the extra curves to facilitate the volume meshing scheme.

If the geometric definition of the surfaces on either side of a given curve is the same the curve can usually be "regularized" away using a solid modeling kernel operation. However, when this is not possible, virtual composite operations are used to composite the surfaces on either side of the curve together into a single surface. Fig. 9 shows the final result after removing the extraneous curves.

4 Advantages of the Small Feature Removal Process

One advantage of this small feature removal process is that most of the operation is done using the solid modeling kernel. As a result, the modifications

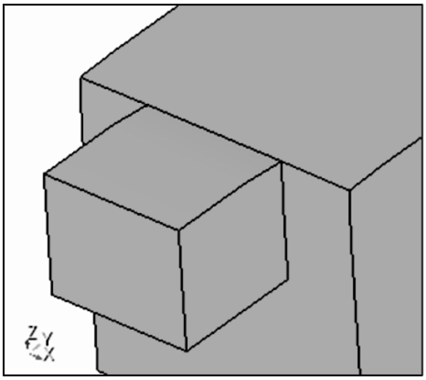

Fig. 9. Final result after removing extraneous curves

can be saved in the native B-rep format and reused in other applications. The only part of the operation that uses virtual topology is the clean-up of the extraneous curves at the end using composites. Obviously, these composites cannot be saved and used in other applications. However, the significant part of the topology/feature removal is persistent and the composites can be reapplied if necessary when needed.

Another advantage of this small feature removal process is that it can remove complex networks of "small" features and results in a fairly smooth transition from the new topology to the rest of the model. If the narrow surfaces in the above example were all eliminated using just composite surfaces, the composite surfaces would contain extreme C1 discontinuities where the surfaces meet at 90 degree angles. These types of discontinuities can greatly inhibit attempts to mesh the composite surface with advancing front surface meshers that rely heavily on surface normals. As the meshers advance across sharp corners like this they often break down.

Another advantage of the remove topology operator is that it could be incorporated into a CAD design system as it uses mainly real solid modeling kernel operations. Even though CAD systems provide powerful features for design they don't provide a large set of tools for removing unwanted topology not related to features.

5 Example

This example is of a part that contains many very small features on the order of 1e-5 but where the desired mesh size is about 0.2. Fig. 10 and Fig. 11 show zoomed in pictures of some of the small features. Fig. 10 (a) shows a cutout where a single curve is expected. The cutout propagates all the way along the top of the model. Fig. 10 (b) shows a cylindrical surface coming in tangent to another surface. The top of the cylindrical surface comes short of meeting

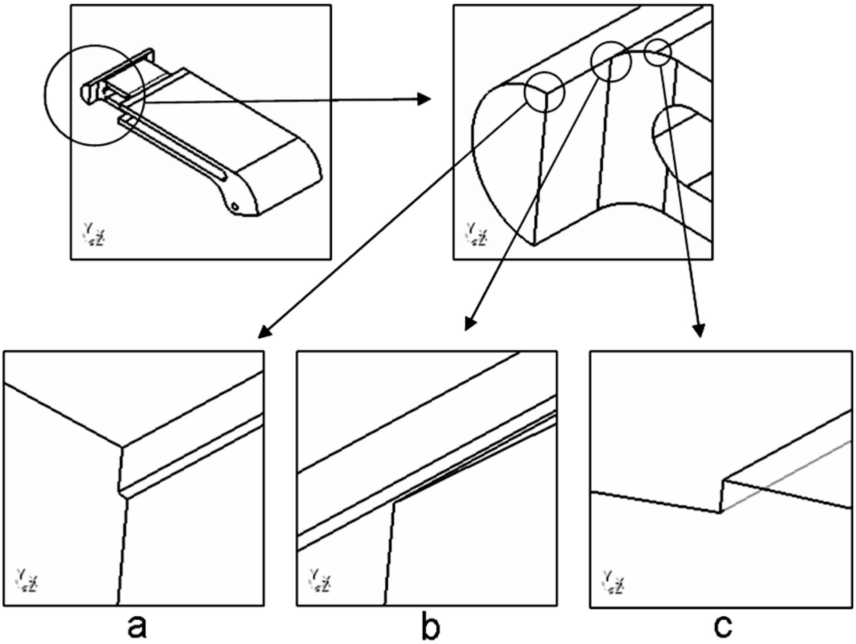

Fig. 10. Close-ups of problems in example model: a) unexpected cutout instead of single curve, b) features not connected as expected, c) unexpected small step

up with the cutout from Fig. 10 (a) creating a very narrow region. Fig. 10 (c) shows a very small step where a single curve is expected. Fig. 11 (a) shows a very narrow region that exists because two features don't line up exactly. Similarly, Fig. 11 (b) shows a step caused by the same misalignment. Fig. 11 (c) shows a small step with another small surface at the base of the step.

A tetrahedral mesher was used to generate a mesh on the unmodified geometry. However, the mesher could not succeed with the tiny features present. The automatic surface splitting algorithm was used to prepare the model for the remove topology operator which was then applied to the small features. Fig. 12 shows one of the small features as it goes through the split surface and remove topology operations. Higher zoom magnifications are shown in the progressing rows and the split surface and remove topology operation results in the progressing columns. Column 2 of Fig. 12 shows the model after applying the surface splitting algorithm. In row 2 (first zoom level) the surface with the curves meeting tangentially is split to remove the portion of the surface that is considered "narrow". The algorithm stepped back from the tangency until the distance between the two curves is the characteristic small curve length. It then split the surface at that location. As can be seen by this example, the process can be used to remove small angles in the model. In row 3 (second zoom level) the narrow region on the surface where the cylinder

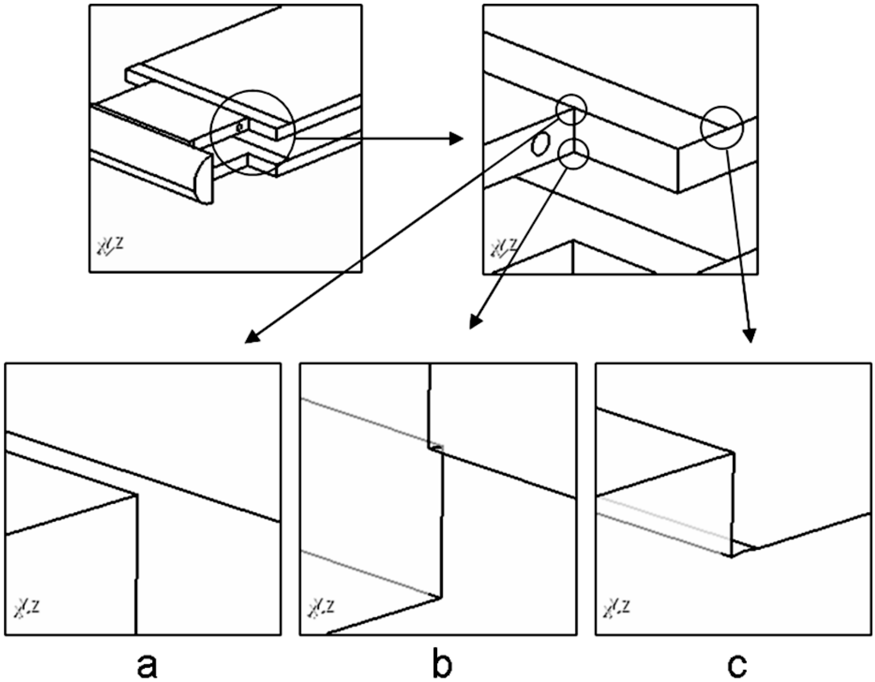

Fig. 11. More problems: a), b) features don't align as expected c) unexpected step and extra small surface

comes in tangent was split out and then the new vertex introduced by this split was propagated to the other adjacent narrow surfaces. All of these new vertices will form a single clump in the new topology.

Column 3 of Fig. 12 shows the model after the remove topology operator was applied. As can be seen, the topology was greatly simplified removing all of the little steps and misalignments that existed previously. The dotted lines indicate extraneous curves that were composited out and show where the original model was cut to remove the old topology Column 4 shows the final topology with the composited curves not drawn. This is the topology that the meshing algorithms interact with.

After applying the small feature removal process to all of the small features in the model the tetrahedral mesher was successful in generating the mesh with size 0.2 as shown in Fig. 13. Note that there is no unnecessary refinement which would have been necessary if the small features were present (the refinement in the far right frame is due to a hole that is part of the design).

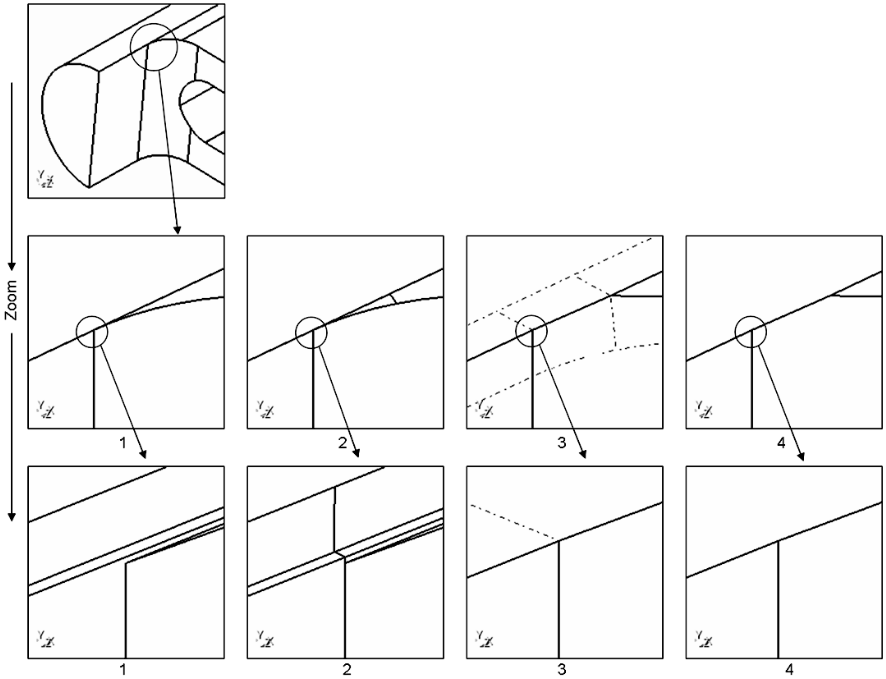

Fig. 12. Zoom in from top to bottom. Split surface and remove topology operations from left to right: 1) unmodified model 2) after split surface algorithm 3) after remove topology (showing extraneous curves as dotted lines) 4) after compositing out extraneous curves

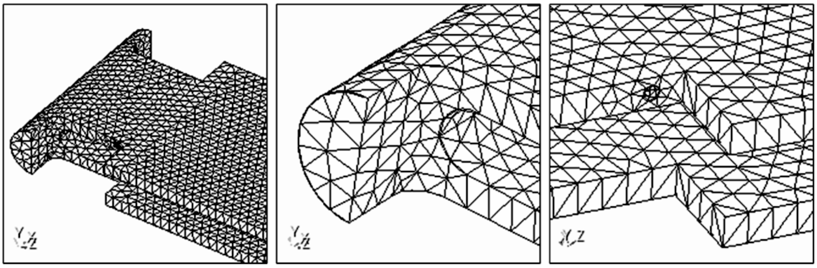

Fig. 13. Meshed example after removing small topology

6 Conclusion and Future Work

This paper introduced a new process for removing small features from a solid model using real operations provided by the solid modeling kernel. The process cuts out the specified topology and reconstructs the topology and geometry required to close the volume and maintain a water-tight topology. The process is built upon functionality that is commonly provided by solid modeling kernels. The process provides a general way to remove unwanted features and provides powerful capabilities for general geometry cleanup. Examples were given that demonstrated the process's ability to remove small curves, small/narrow surfaces, small regions in surfaces, and complex combinations of these. It was also shown to be able to remove small angles created by surface tangencies in the B-rep model.

One area for future research is to automate the application of this small feature removal process to a model in a global fashion. The algorithm relies on the user to provide the topology to be removed, a small curve length to decide what is "small", and the distance to backoff from the topology being removed when cutting it out of the model. These values will vary at different locations in the model and for different features to be removed. Having an algorithm that automatically generated and specified the input to the process based on the analysis of the B-rep model would be valuable.

A second area for future research is to develop a new method or incorporate an existing method for reconstructing the surfaces adjacent to the topology being removed so that they maintain their original shape as much as possible and also connect to the new topology. This capability could be used in 2 different ways: 1) the current algorithm could use this capability to reconstruct a single surface representing surfaces that are currently being composited together using virtual topology or 2) the current algorithm could be simplified to not "cut out" the old topology but rather just reconstruct surfaces adjacent to the old topology so that they connect to the new topology. This sort of reconstruction functionality would have an effect similar to a virtual composite operation but would be provided using real solid modeling operations.

A third area for future research is to extend the remove topology operator to work on assemblies. Adjacent volumes in an assembly often contain topology (vertices, curves, and surfaces) that occupies the same space within some tolerance. This topology is often "merged" together into a single piece of topology representing the two. The surrounding topology of both volumes is modified to incorporate the merged topology. This merging is often done in order to generate a contiguous mesh between the two volumes. If the topology being removed during a remove topology operation is topology that could be merged with topology on adjacent volumes, the remove operation should be performed on all such volumes at the same time so that the resulting topology can also be merged once the operation is complete. As well as maintaining mergeable topology this could be used as a tool to fix misalignments between volumes that prevent correct merging.

7 Acknowledgements

Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

References

1. T.D. Blacker. *CUBIT Mesh Generation Environment, V. 1: User's Manual.* Sandia National Labs, 1994. SAND94-1100.
2. Spatial Corporation. Acis 3d toolkit. http://www.spatial.com/products/acis.html.
3. N.C. Eccles and J.P. Steinbrenner. Solid modeling and fault tolerant meshing-two complimentary strategies. In *17th AIAA Computational Fluid Dynamics Conference*, June 6-9, 2005.
4. G. Butlin and C. Stops. Cad data repair. In *Proceedings of 5th International Meshing Roundtable*, 1996.
5. M.R. Jones, M.A. Price, and G. Butlin. Geometry management support for auto-meshing. In *Proceedings of 4th International Meshing Roundtable*, pages 153–164. Sandia National Laboratories, 1995.
6. K.Y. Lee, C.G. Armstrong, M.A. Price, and J.H. Lamont. A small feature suppression/unsuppression system for preparing b-rep models for analysis. In *Proceedings 2005 ACM Symposium on Solid and Physical Modeling*, pages 113–124, June 13-15, 2005.
7. S. Venkataraman and M. Sohoni. Reconstruction of feature volumes and feature suppression. In *Proceedings 7th ACM Symposium on Solid Modeling and Applications*, pages 60–71, 2002.
8. A. Sheffer, T.D. Blacker, J. Clements, and M. Bercovier. Virtual topology construction and applications. In *Proceedings Theory and Practice of Geometric Modeling*, 1996.
9. T.J. Tautges. Automatic detail reduction for mesh generation applications. In *Proceedings 10th International Meshing Roundtable*, pages 407–418, 2001.
10. A. Sheffer. Model simplification for meshing using face clustering. *Computer Aided Design*, 33:925–934, 2001.
11. C.G. Armstrong, S.J. Bridgett, R.J. Donaghy, R.W. McCune, R.M. McKeag, and D.J. Robinson. Techniques for interactive and automatic idealisation of cad models. In *Proceedings 6th International Conference on Numerical Grid Generation in Computational Field Simulations*, pages 643–662, 2001.
12. A.V. Mobley, M.P. Carroll, and S.A. Canann. An object oriented approach to geometry defeaturing for finite element meshing. In *Proceedings 7th International Meshing Roundtable*. Sandia National Labs, 1998.
13. S. Dey, M.S. Shephard, and M.K. Georges. Elimination of the adverse effects of small model features by the local modification of automatically generated meshes. *Engineering with Computers*, 13:134–152, 1997.
14. M.S. Shephard, M.W. Beall, and R.M. O'Bara. Revisiting the elimination of the adverse effects of small model features in automatically generated meshes. In *Proceedings 7th International Meshing Roundtable*, pages 119–131. Sandia National Labs, 1998.