Mason: Morphological Simplification

Jason Williams[1] and Jarek Rossignac

College of Computing, IRIS cluster, GVU Center

Georgia Institute of Technology, Atlanta GA, USA

http://www.gvu.gatech.edu/~jarek/iris/

ABSTRACT

Morphological filters, such as closure, opening, and their combinations, may be used for cleaning and analyzing images and shapes. We focus on the most popular special cases of these operators: the rounding $R(S)$ and the filleting $F(S)$ of an arbitrary set S and the combinations $R(F(S))$ and $F(R(S))$. These operators may be obtained by combining growing and shrinking operators, which are Minkowski sums and differences with a ball of a given radius r. We define the mortar $M(S)$ as $F(S) \setminus R(S)$. Note that the mortar occupies the thin cracks, protrusions, constrictions, and areas near the high curvature portions of the boundary of S. Thus, we argue that confining the effect of shape simplification to the mortar has advantages over previously proposed tolerance zones and error metrics, which fail to differentiate between the irregular regions contained in the mortar and the regular (low-curvature) regions of S. We point out that $R(F(S))$ and $F(R(S))$ are suitable simplification filters in this context, because their effects are confined to $M(S)$ and leave the core $R(S)$ and the anticore $\overline{F(S)}$ unchanged. Furthermore, they tend to replace the high-curvature portions of the boundary of S with regular portions where the radius of curvature exceeds r. Unfortunately, these operators have a bias, which may result in a large total volume of the symmetric difference between S and its simplified version S'. In order to minimize this volume, we propose to select the filter locally, for each connected component of the mortar. Thus, some portions of the mortar will be simplified using $F(R(S))$ and some using $R(F(S))$. This approach, which we call the Mason filter, can be used for the simplification of shapes regardless of their representation or

[1] Corresponding author. Address: 426 Marietta St. NW Apt. 410 Atlanta, GA 30313
Phone: (404) 688-9713 Email: jasonw@cc.gatech.edu

dimensionality. We demonstrate its application to discrete two-dimensional binary sets (i.e. black and white images) and discuss implementation details.

Keywords: mathematical morphology, shape simplification, binary image processing

1. INTRODUCTION

The objective of this work is to explore new shape simplification techniques that are suitable for multiresolution visualization and analysis. Our long-term goal is to support analysis of how a model's geometry and topology change as we progressively simplify it. We seek to simplify the high frequency (detailed) features of a shape S by replacing them with more regular features, but to leave unchanged the portions of the shape and of its complement that lie away from these features.

Prior simplification solutions lacked a formal characterization of the regions with high frequency features and thus were not able to confine the simplification to these regions. To overcome this deficiency, we provide a formal definition of the region that surrounds high frequency features. We call it the *mortar* of S and denote it $M(S)$. Thus, we require that the filters do not alter the *core* of S, defined as $S \ominus M(S)$, and the *anticore* of S defined as $\overline{S} \ominus M(S)$. Figure 1 shows a shape S (black) on the left. On the right, we distinguish its core (black), its mortar (blue), and its anticore (white).

Figure 1: *A shape (black) on the left, and its core (black), mortar (blue), and anticore (white) on the right.*

We propose a new morphological filter, which we call **Mason**, that **selectively carves portions of the mortar**, attempting to create a shape that has few fine details while **minimizing the cost**, which we define as the area (in 2D) or the volume (in 3D) of the portion of space affected by the change. The Mason filter can be used to simplify shapes regardless of their representation. We show the results it produces on 2D black and white (binary) images and compare them to the results of applying combinations of rounding and filleting.

In the remainder of the paper, we review prior art in shape simplification, provide definitions and properties used to define the Mason approach, and present the Mason approach, implementation details, and results.

2. PRIOR ART

We cannot draw on many of the previously proposed simplification methods, because they do not support our goal of analyzing how topology changes when we selectively remove features below a given scale. For instance, we cannot use simplification techniques that preserve topology, such as most methods that incrementally reduce the number of triangles in a mesh [CIG98a]. Of methods that allow topology changes, most modify a model's geometry as an unwanted side effect of reducing the size of its representation. They are not designed to remove all detail below a given scale, and high degrees of simplification may even introduce details such as sharp angles. In addition, most error metrics do not differentiate between detailed and undetailed regions in the model, so they cannot confine changes to detailed regions. These observations hold for diverse simplification techniques, including topology-modifying triangle reduction [POP97], octrees [AND02], wavelets [CHO97], and vertex clustering [ROS93], as well as diverse approaches to tolerance zones and error metrics, including simplification envelopes [CHO96] and minimization of the Hausdorff error [CIG98b], maximum error [RON96], and a quadratic error [GAR97].

The approach in [WOO] is notable in that it guarantees removal of all handles below a given scale. However, it does not remove geometric features or reduce the number of connected components. [VAN03] removes all handles, regardless of scale, and reduces the model to a single connected component. It also preserves most geometric features.

[HE96] applies a low-pass filter in the volume domain. Although this approach offers some control over the scale of the features that are simplified away, it tends to erode the model, even along regular portions. By contrast, the morphological operators upon which our approach is built do not displace regular contours.

Although the approach of [ELS98] is framed in terms of alpha hulls and operates on polygonal meshes, it achieves results comparable to rounding and filleting. [NOO03] use standard rounding and filleting to simplify the topology of volumetric data obtained from polygonal models.

A discussion of prior art in morphological filters is deferred to the next section.

3. DEFINITIONS AND PROPERTIES

3.1. Minkowski operators

Our method expands on basic mathematical morphology operators. Classic texts on the subject include [SER82] and [HEI94]. Although mathematical morphology generalizes to lattices (partially ordered sets whose subsets each have a greatest lower bound and a least upper bound,) we restrict ourselves to regularly sampled sets and discuss implementation results for sets of pixels of a 2D black and white image.

Many operations in mathematical morphology can be defined in terms of the Minkowski sum and difference. The ***Minkowski sum*** of A and B, denoted $A \oplus B$, is $\bigcup_{b \in B} A + b$; it is the result of taking the union of copies of B (which is called the structuring element) translated to every point in A. The ***Minkowski difference***, denoted $A \ominus B$ is $\bigcap_{b \in B} A - b$; it is the set of points p such that B translated to p lies completely within A. An alternative definition is $A \ominus B = \overline{\overline{A} \oplus \hat{B}}$, where $\hat{B}$ is B reflected about the origin.

The Minkowski sum and difference are not inverses. $(A \ominus B) \oplus B$, for instance, is the union of all translated copies of B that fit in A. For fixed B, we can write $(A \ominus B) \oplus B$ as an operator $\gamma(A)$. $\gamma(A)$ is an ***opening***, which is defined as an operation that is translation-invariant $(\gamma(A + b) = \gamma(A) + b)$, idempotent $\gamma(\gamma(A)) = \gamma(A)$, increasing $(\gamma(A) \subseteq \gamma(C) \subseteq A \subseteq C)$, and antiextensive $(\gamma(A) \subseteq A)$. The dual operator $\phi(A) = (A \oplus \hat{B}) \ominus \hat{B}$ is the complement of all the copies of B that fit in $\overline{A}$. It is a ***closing***, which is defined as an operation that is translation-invariant, idempotent, increasing, and extensive $(\phi(A) \supseteq A)$.

3.2. Growing, shrinking, rounding, and filleting operators

We are particularly interested in the case where the structuring element B is a ball b_r of radius r. Thus, we define S ***grown*** by r to be $S\uparrow_r = S \oplus b_r$, S ***shrunk*** by r to be $S\downarrow_r = S \ominus b_r$, S ***filleted*** by r to be $F_r(S) = S\uparrow_r\downarrow_r$ and S ***rounded*** by r to be $R_r(S) = S\downarrow_r\uparrow_r$ [ROS84] [ROS85] [ROS86] . (We will omit the subscript *r* when dealing with only a fixed radius of simplification.) $S\uparrow$ is the union of all balls whose centers lie in S, $S\downarrow$ is the centers of all balls that lie in S, $R(S)$ is the union of all balls that lie in S, and $F(S)$ is the complement of the union of all balls that lie in $\overline{S}$, or equivalently, the region inaccessible to a ball outside of S

(Figure 2). $R(S)$ is an opening, and so it is antiextensive, while $F(S)$ is a closing, so it is extensive. From this we obtain that $R(S) \subseteq S \subseteq F(S)$.

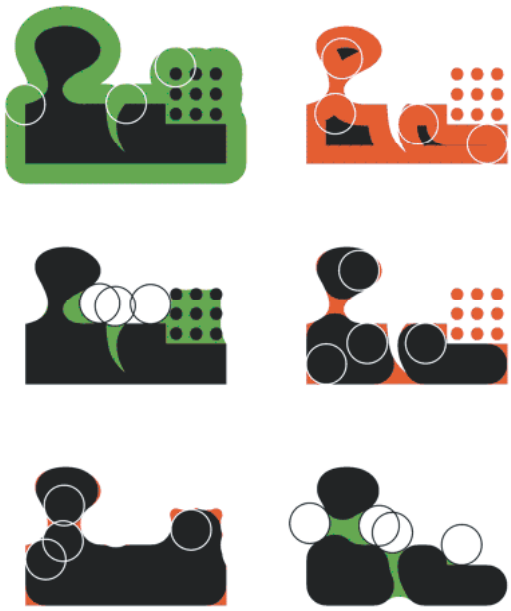

Figure 2: *Morphological operations applied to the set S from Figure 1. The material the filters add is shown in green, while the material they remove is shown in red. (Top left)* $S\uparrow$ *(Top right)* $S\square$ *(Middle left)* $F(S)$ *(Middle right)* $R(S)$ *(Bottom left)* $R(F(S))$ *(Bottom right)* $F(R(S))$.

3.3. Mortar, core, and anticore

We have defined the ***mortar*** $M(S)$ as $F(S) \setminus R(S)$. Note that the mortar is concentrated around the boundary of S in regions of high curvature and in thin regions between two distinct portions of the boundary (Figure 1). We have defined he ***core*** as $S \setminus M(S)$. The core is the thick portion of S, which has no sharp concave corners, thin branches, or small isolated components. The core will not be affected by simplification. The ***anticore***, defined as $\overline{S} \setminus M(S)$, is the thick portion of $\overline{S}$, which excludes concave corners, thin gaps, or small holes. The anitcore is also not affected by simplification. A symbolic diagram of the set-inclusion relationships between S, $F(S)$, $R(S)$, the core, the mortar, and the anticore appears in Figure 3.

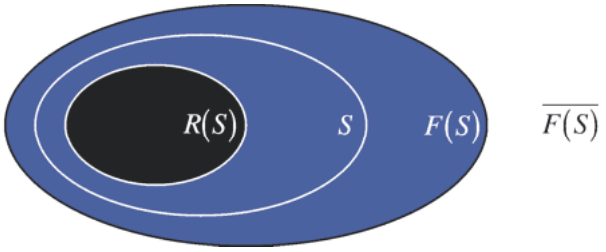

Figure 3: $R(S) \subseteq S \subseteq F(S)$ *The core is shown in black, the mortar in blue, and the anticore in white.*

It leads to a simple proof (see Appendix) of the following two theorems.

Theorem 1: The core, originally defined as $S \circ M(S)$, is $R(S)$.

Theorem 2: The anticore, originally defined as $\overline{S} \circ M(S)$, is $\overline{F(S)}$.

3.4. Inner-regular and outer-regular points

We can now relate our definitions of the core, mortar, and anticore to the definition of r**-regularity** used in [SER82, ATT97]. For this, we define a point in S as ***inner*** r***-regular*** if it is contained in a ball of radius r that lies completely within S, that is, if it is contained in the core $R(S)$. We define a set as inner r-regular if all of its points are inner r-regular, that is, if $S = R(S)$. Similarly we define a point in $\overline{S}$ as ***outer*** r***-regular*** if it is contained in a ball of radius r that lies completely in $\overline{S}$, so that it is contained in the anticore. A set is outer regular if $S = F(S)$. A set is regular if it is both inner regular and outer regular. In this case, $R(S) = S = F(S)$.

The mortar is the set of points that are neither inner-regular or nor outer-regular. Thus it is the union of the inner-irregular set $S \setminus R(S)$, which includes the branches, convex corners, and small components of *S*, with the outer-irregular set $F(S) \setminus S$, which contains the gaps, concave corners, and holes. If S is regular, $M(S)$ is empty.

Because rounding and filleting are idempotent, $R(S)$ and $R(F(S))$ are guaranteed to be inner-regular, while $F(S)$ and $F(R(S))$ are guaranteed to be outer-regular. In practice, $R(F(S))$ and $F(R(S))$ are regular almost everywhere. However, Figure 4 shows a set S for which neither $F(R(S))$ nor $R(F(S))$ is regular. Despite their limitations, a combination of rounding and filleting is among the most effective ways to eliminate (round-off) irregular features.

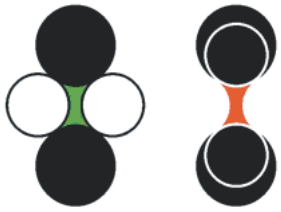

Figure 4: *(Left)* $S = R(S) = R(F(S))$ *is inner-regular but not outer-regular. The outer-irregular points that filleting would add are shown in green. (Right)* $F(S) = F(R(S))$ *is outer-regular but not inner-regular. The inner-irregular points that rounding would remove are shown in red.*

3.5. Bias of R(F(S)) and F(R(S)) filters

The *R(F(S))* and *F(R(S))* filters remove (i.e., round-off) the irregular portions of S, eliminating small holes and merging isolated components. Furthermore, they leave the regular portion unchanged. Hence, they are excellent candidates for shape simplification. Unfortunately, as illustrated in Section 6, each one of these two operators has a bias when operating on portions of the mortar that combine outer-irregular and inner-irregular regions: they either accept them all or reject them all, thus creating a significant change in the area of S. For instance, consider Figure 2. Note that $R(F(S))$ fills in the region containing the small circles, while $F(R(S))$ clears it out. The first operation applied determines whether detail is filled or cleared; once $F(S)$ fills in a region, rounding it only regularizes the region's boundary. Consequently $R(F(S))$ is biased to have more area than S, while $F(R(S))$ is biased to have less.

We would like to eliminate this bias and choose between filling and clearing a region based on which causes the least amount of change to the model. We measure the change in terms of the area affected by simplification. Formally, if S' is the result of applying a particular filter to S, we want to minimize the cost of this simplification, which we define as the area of the symmetric difference between S and S'. For binary images, this amounts to counting the number of pixels that change color. Because which filter causes the least change may vary at different locations in the image, we want the flexibility of choosing between the two filters on a region-by-region basis.

3.6. Local decisions

We cannot make the choice of the filter for each individual point. Therefore, our strategy is to segment the space into a suitable set of regions. For each region, we compute the number of pixels that would change status if we were to use *R(F(S))* and the number of pixels that would change status if we were to use *F(R(S))*. We then replace the region with the result produced by the one of these two filters that corresponds to the smaller count of altered pixels.Unfortunately, most space segmentations will usually lead to sharp irregularities at the inter-region boundaries. For example, a regular lattice of 16x16 pixel regions may result in checkerboard patterns inside the mortar. In the following section, we explain how we identify more suitable candidate regions that avoid most irregularities.

3.7. Regions as connected components of the mortar

The following two theorems are proven in the Appendix.

Theorem 3: $R(S) \cap R(F(S)) \cap F(S)$

Theorem 4: $R(S) \cap F(R(S)) \cap F(S)$

Combining them, we can conclude the following

Theorem 5: The result of applying a combination of rounding and filleting to S is guaranteed to differ from S only in the mortar.

Thus we can use the maximally connected components, M_i, of the mortar as regions and replace each one of them with either $M_i \cap F(R(S))$ or $M_i \cap R(F(S))$, selecting the one with the smaller cost (i.e. count of altered pixels). This idea leads to a simple and effective algorithm, which we call Mason, discussed in details in the next section.

4. IMPLEMENTATION DETAILS

4.1. Identifying the mortar

We first compute $M(S)$, $R(F(S))$, and $F(R(S))$ by combining the results produced by the standard $F(S)$ and $R(S)$ filters. Researchers have proposed several algorithms for computing morphological operations, which

vary depending on the representation of the operands and the nature of the structuring element. [VAR], for instance, addresses the problem of approximately computing the Minkowski sum of two polyhedral models. [ROS84], [ROS85], and [ROS86] present methods for computing morphological operations involving CSG models and a ball, while [ROD03] details an efficient algorithm for computing operations between volume data and a box-shaped structuring element. [CUI99], which is more relevant to our problem, describes a method for computing morphological operations between volume data and a ball by computing a distance transform through progressively expanding the volume's boundary. We recommend use of this method when the radius of simplification is large. However, we opt for a simpler implementation that is efficient for small radii, perhaps up to 10 pixels.

To compute $F(S)$, we scan through the image with a mask containing all pixels within a distance r of the center pixel. If when the mask is translated to a pixel p all pixels covered by the mask are white, we mark all the pixels in the mask. As a result of such a pass, marked pixels are those fitting in a ball of radius r contained in $\overline{S}$. The unmarked pixels belong to $F(S)$. We use a similar procedure to compute $R(S)$, marking the pixels belonging to $R(S)$ when all the pixels within the mask are black.

4.1. Identifying the regions

To identify the connected components of the mortar, we associate a label with each pixel. We assign pixels in the core and anticore the same label, while pixels in the mortar are initially unlabeled. We scan through the image until we reach an unlabeled mortar pixel. We record the pixel in a list with a new label. We then start a breadth-first traversal by labeling the pixel and its four-connected neighbors and adding the neighbors to a queue. Then, until the queue is empty, we de-queue a pixel, label its unlabeled neighbors, and add them to the queue. When this process completes, all the mortar pixels in the same connected component share the same label. We continue the scan and repeat the process with a new label at the next unlabeled pixel. The size of the stack may become prohibitive for large models. We have also implemented a more efficient traversal of the mortar components that was inspired by the EdgeBreaker [ROS99] compression traversal for triangle meshes and by its extension to quadrilateral meshes [KIN99]. During the traversal, each pixel, treated as a quad, is split into two triangles. They are invaded one at a time. When the invasion of a triangle does not split the mortar

component, there is no need for recursion. Recursion is needed only in the very rare cases when all three vertices of the new triangle have already been removed from the mortar and when exactly two of its edge-adjacent triangles are still in the mortar.

4.2. Selecting the best filter for each region

For each connected component M_i of the mortar, we traverse the pixels in M_i and compute the count *fr* of pixels of M_i where *S* and *F(R(S))* disagree and the count *rf* of pixels of M_i where *S* and *R(F(S))* disagree. If *fr*<*rf*, then we copy the pixels from $F(R(S)) \cap M_i$ to our output image $Mason(S)$; otherwise, we copy the pixels from $R(F(S)) \cap M_i$.

5. RESULTS

We present the images produced by applying combinations of rounding and filleting, as well as the Mason filter, to an image of a landscape (Figure 5.)

We provide statistics indicating for each image: the area, the number of pixels changed, the number of black components, and the number of white components. We also measure the number of inner-irregular and outer-irregular pixels (Table 1.) The images are 400 by 300 pixels and were simplified with a radius of 5. The inner irregular pixels are shown in green, while the outer irregular pixels are shown in red.

As the results indicate, $F(R(S))$, $R(F(S))$, and $Mason(S)$ yield similar, and substantial, reductions in the topological complexity and irregularity of S. $Mason(S)$ is neither inner regular nor outer regular, but it has a smaller number of pixels changed and a smaller change in total area than the other filters. As noted earlier, $F(R(S))$ fills in thick connected components of the mortar, while $R(F(S))$ clears them out. $Mason(S)$ chooses between filling and clearing based on which results in the smallest number of pixels changed. For instance, Mason chooses to use $F(R(S))$ to paint white the predominantly white portion of water to the right of the tree and chooses to use $R(F(S))$ to paint black the predominantly black portion of the top left part of the image. In this image, some components contain mostly inner irregular pixels while others contain mostly outer

irregular pixels, so Mason's pattern of filling and clearing achieves significantly lower cost than either filling in all the components or clearing them all out.

6. LIMITATIONS

Figure 6 illustrates such a situation where Mason's advantages are less pronounced. In this 300 by 225 pixel image of a slide of lung cells, simplified with a radius of 5, most of the connected components of the mortar contain comparable quantities of inner irregular and outer irregular pixels, so all patterns of filling and clearing yield similar numbers of pixels changed.

Similarly, Figure 7 and Table 3 highlight a limitation shared by $Mason(S)$, $F(R(S))$, and $R(F(S))$. In this 620 by 430 pixel image of a galaxy, also simplified with a radius of 5, predominantly white regions with black noise are adjacent to predominantly black regions with white noise, forming a single connected component of mortar. In this case, Mason is forced to make a single decision for the mixed component and hence loses its advantage over $R(F(S))$ and $F(R(S))$. $F(R(S))$ clears out this component, while $R(F(S))$ and $Mason(S)$ fill it in. All three filters result in a large area change.

Figure 5: *(Top left) The original image* S*. (Top right) The inner irregular points of* S *are colored red, and the outer irregular points are colored green. The core is black, while the anticore is white. (Middle left)* $F(R(S))$ *Note the small bridges of inner irregular points marked in red.(Middle right)* $R(F(S))$ *Note the green bridge of outer irregular points in the upper right portion of the figure. (Bottom left)* $Mason(S)$.

	Area	Pixels Changed	Black Components	White Components	Inner-irregular Pixels	Outer-irregular Pixels
S	78,195	N/A	527	711	18,676	18,641
$F(R(S))$	62,808	19,301	6	23	107	0
$R(F(S))$	95,472	18,573	1	33	0	43
$Mason(S)$	75,674	15,071	3	27	48	13

Table 1: Quantitative results for the images in Figure 5.

	Area	Pixels Changed	Black Components	White Components	Inner-irregular Pixels	Outer-irregular Pixels
S	36,967	N/A	98	177	7,970	7,312
$F(R(S))$	31,873	9,218	8	17	203	0
$R(F(S))$	41,842	8,523	5	20	0	89
$Mason(S)$	38,013	7,922	5	22	203	108

Table 2: Quantitative results for the images in Figure 6.

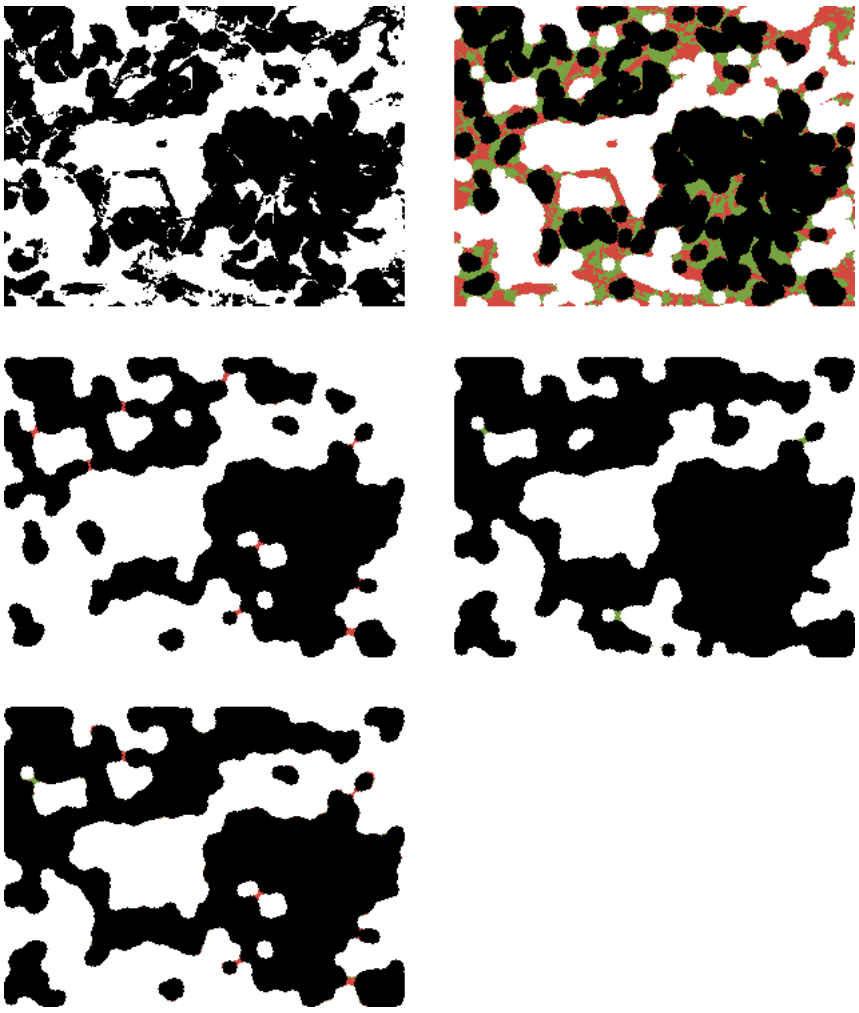

Figure 6: *(Top left) The original image* S. *(Top right) The inner irregular points of* S *colored red, and the outer irregular points colored green (Middle left)* $F(R(S))$ *(Middle right)* $R(F(S))$ *(Bottom left)* $Mason(S)$

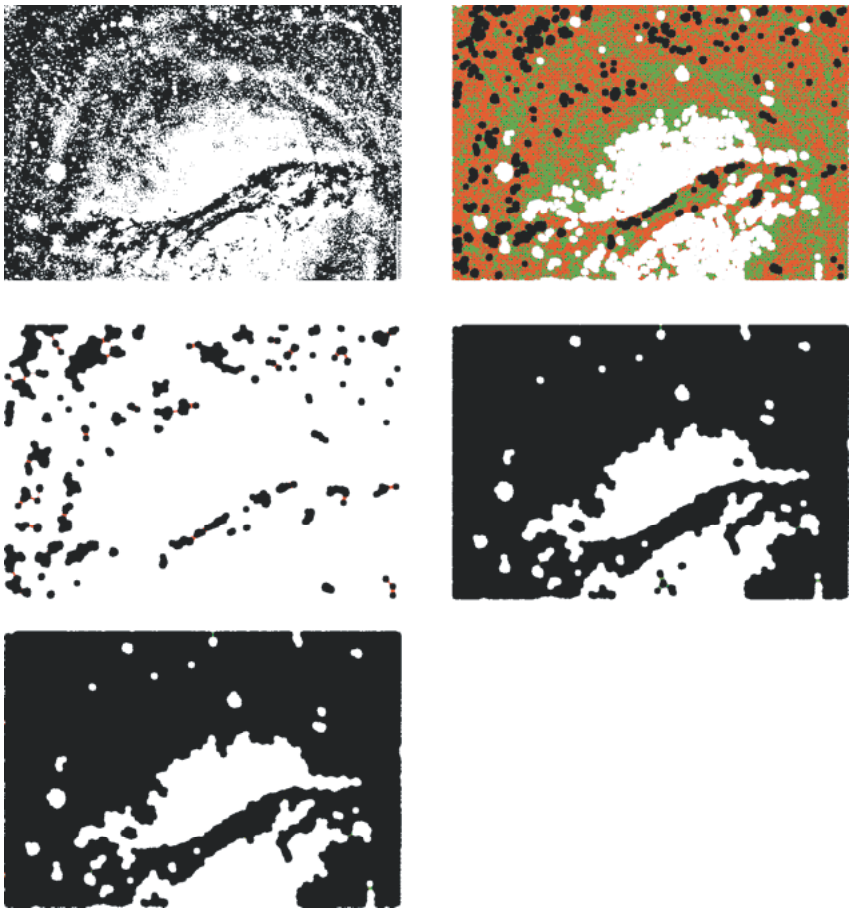

Figure 7: (Top left) Original image S. (Top right) The inner-irregular points of S colored red, and the outer irregular points colored green. (Middle left) $F(R(S))$ (Middle right) $R(F(S))$ (Bottom left) $Mason(S)$

	Area	Pixels Changed	Black Components	White Components	Inner-irregular Pixels	Outer-irregular Pixels
S	128,956	N/A	2,212	3,218	101,717	79,769
$F(R(S))$	29,993	100,969	64	6	679	0
$R(F(S))$	204,213	79,225	9	145	0	154
$Mason(S)$	203,329	78,923	3	145	15	107

Table 3: Quantitative results for the images in Figure 7.

7. DISCUSSION

Given a shape S, our problem has been to find a nearly r-regular shape S' such that the symmetric difference $S \sqcap S'$ has minimal area and lies completely within the mortar $M(S)$. Ideally our method would produce a fully r-regular shape. However, in some cases, such as in Figure 2, there is no r-regular S' such that $S \sqsubseteq S' \sqsubseteq M(S)$.

The filters $R(F(S))$ and $F(R(S))$ are dual of each other, meaning that $R(F(S)) = \overline{F(R(\overline{S}))}$. The mason filter, by contrast, is self-dual. If we write the result of applying it as $Mason(S)$, we have that $Mason(S) = \overline{Mason(\overline{S})}$. Unlike rounding and filleting, the Mason filter treats positive and negative space symmetrically. In this respect it achieves our goal of eliminating the bias in $R(F(S))$ and $F(R(S))$.

8. SUMMARY OF CONTRIBUTIONS AND CONCLUSIONS

Previously proposed shape simplification techniques used Hausdorff or other error measures for shape simplification that allowed the shape to be altered both in regular and in irregular regions. We have proposed a new tolerance zone, called the mortar, that confines the effects of shape simplification to irregular regions, where the boundary has high curvature or where two distinct portions of the boundary are close to each other. Such a confinement ensures that the low-curvature, regular portions of the boundary are not affected by simplification. We have developed a mathematically simple set-theoretical definition of the mortar that is independent of the dimension of the underlying space and of the choice of the representation. Furthermore, we have demonstrated that the mortar may be identified using a simple and efficient implementation.We have proven that standard mathematical morphology operators (combinations of rounding and filleting) only alter the shape inside the mortar, and we have shown that they tend to produce shapes that are regular almost everywhere.We have pointed out that for some portions of the image, one of these operators may be better, while the other is more suitable for other portions. We have therefore proposed to split the image into regions and to select the most suitable filter for each region independently. To ensure that the disparity of filter choices between regions does not result in unexpected irregularities, we have devised an approach where the regions are identified as the connected component of the mortar. We measure the suitability of a filter as the count of pixels

it changes. We have demonstrated that our approach, which we call the Mason filter, produces nearly r-regular shapes with less change to the model than combinations of standard rounding and filleting operations.

9. ACKNOWLEDGEMENTS

This research was supported by a DARPA/NSF CARGO grant number 0138420.

REFERENCES

[AND02] C. Andujar, P. Brunet, and D. Ayala. Topology-reducing surface simplification using a discrete solid representation. ACM Transactions on Graphics 21, 2, 88-105. 2002.

[ATT97] D. Attali. r-Regular shape reconstruction from unorganized points. Computational Geometry, 248-253. 1997.

[CHO97] M. Chow and M. Teichmann. A wavelet-based multiresolution polyhedral object representation. Proceedings of ACM SIGGRAPH '97, technical sketch. 1997.

[CIG98a] P. Cignoni, C. Montani, and R. Scopigno. A comparison of mesh simplification algorithms. Computers and Graphics 22, 1, 37-54. 1998.

[CIG98b] P. Cignoni, C. Rocchini, and R. Scopigno. Metro: measuring error on simplified surfaces. Computer Graphics Forum 17, 2, 167-174. 1998.

[COH96] J. Cohen, A. Varshney, et. al. Simplification envelopes. Proceedings of ACM SIGGRAPH, 119-128. 1996.

[CUI99] O. Cuisenaire. Distance transformations: fast algorithms and applications to medical image processing. Ph.D. thesis, Universite catholique de Louvain. 1999.

[ELS98] J. El-Sana and A. Varshney. Topology simplification for polygonal virtual environments. IEEE Transactions on Visualization and Computer Graphics 4, 2, 133-144. 1998.

[FAI03] S. Fairy. Available at http://www.cytology-asc.com/branch/nsw/jun03/cases.html. 2003.

[GAR97] M. Garland and P. Heckbert. Surface simplification using quadric error metrics. Proceedings of ACM SIGGRAPH.1997.

[HE96] T. He, L. Hong, et. al. Controlled topology simplification. IEEE Transactions on Visualization and Computer Graphics 2, 2, 171-183. 1996.

[HEI94] H. Heijmans. Morphological Image Operators. Academic Press. 1994.

[KIN99] D. King and J. Rossignac. Connectivity Compression for Irregular Quadrilateral Meshes. Georgia Institute of

Technology, GVU Technical Report GIT-GVU-99-36.

[NOO03] F. Nooruddin and G. Turk. Simplification and repair of polygonal models using volumetric techniques. IEEE Transactions on Visualization and Computer Graphics 9, 2, 191-205. 2003.

[POP97] J. Popovic and H. Hoppe. Progressive simplicial complexes. Proceedings of ACM SIGGRAPH, 217-224. 1997.

[ROD03] J. Rodriguez, D. Ayala, and A. Aguilera. A complete solid model for surface rendering. Universitat Politecnica de Catalunya, Report LSI-03-3-R. 2003.

[RON96] R. Ronfard and J. Rossignac. Full-range approximations of triangulated polyhedra. Computer Graphics Forum 15, 3, C-67. 1996.

[ROS85] J. Rossignac. Blending and offsetting solid models. Ph.D. thesis, University of Rochester, NY. 1985.

[ROS99] J. Rossignac. Edgebreaker: Connectivity compression for triangle meshes. IEEE Transactions on Visualization and Computer Graphics, 5, 1, 47-61. 1999.

[ROS93] J. Rossignac and P. Borrel. Multi-resolution 3D approximations for rendering complex scenes. In Geometric Modeling in Computer Graphics, ed. B. Faliciendo and T. Kunii. Springer-Verlag, 455-465. 1993.

[ROS84] J. Rossignac and A. Requicha. Constant-radius blending in solid modeling. ASME Computers in Mechanical Engineering, 3, 65-73. 1984.

[ROS86] J. Rossignac and A. Requicha. Offsetting operations in solid modelling. Computer-Aided Geometric Design 3, 129-148. 1986.

[SER82] J. Serra. Image Analysis and Mathematical Morphology. Academic Press. 1982.

[STE02] R. Stenger. Galactic collision leaves glowing scar. Available at http://www.cnn.com/2002/TECH/space/10/25/galactic.crash. 2002.

[VAR] G. Varadhan and D. Manocha. Minkowsi sum approximation of complex polyhedral models. Available at http://gama.cs.unc.edu/mink.

[WOO] Z. Wood, H. Hoppe, et. al. An out-of-core algorithm for isosurface topology simplification. Submitted for publication. Available at http://www.multires.caltech.edu/pubs/topo_filt.pdf.

APPENDIX

Theorem 1: The core is $R(S)$.

Proof: By definition, the core is $S \cap M(S)$ and $M(S) = F(S) \cap R(S)$. Rewriting $A \cap B$ as $A \cap \overline{B}$, the core is then $S \cap \overline{F(S)} \cap \overline{R(S)}$. Because $\overline{A \cap B} = \overline{A} \cap \overline{B}$, this is equal to $S \cap \left(\overline{F(S)} \cap R(S)\right)$. Distributing the intersection yields $\left(S \cap \overline{F(S)}\right) \cap \left(S \cap R(S)\right)$. Because $S \subseteq F(S)$, $S \cap \overline{F(S)}$ is empty, and because $R(S) \subseteq S$, $S \cap R(S) = R(S)$. Therefore the core is $R(S)$.

Theorem 2: The anticore is $\overline{F(S)}$.

Proof: The proof is similar to the proof of Theorem 1. By definition, the anticore is $\overline{S} \cap M(S)$. This is equal to $\overline{S} \cap \overline{F(S)} \cap \overline{R(S)}$, which equals $\overline{S} \cap \left(\overline{F(S)} \cap R(S)\right)$. Distributing the intersection yields $\left(\overline{S} \cap \overline{F(S)}\right) \cap \left(\overline{S} \cap R(S)\right)$. Because $S \subseteq F(S)$, $\overline{F(S)} \subseteq \overline{S}$, so $\overline{S} \cap \overline{F(S)} = \overline{F(S)}$. Because $R(S) \subseteq S$, $\overline{S} \cap R(S)$ is empty. Therefore the anticore is $\overline{F(S)}$.

Theorem 3: $R(S) \subseteq R(F(S)) \subseteq F(S)$

Proof: Because $S \subseteq F(S)$ and $A \subseteq B \subseteq R(A) \subseteq R(B)$, $R(S) \subseteq R(F(S))$. Because $R(A) \subseteq A$, $R(F(S)) \subseteq F(S)$. Therefore $R(S) \subseteq R(F(S)) \subseteq F(S)$.

Theorem 4: $R(S) \subseteq F(R(S)) \subseteq F(S)$

Proof: The proof is similar to the proof of Theorem 3. Because $A \subseteq F(A)$, $R(S) \subseteq F(R(S))$. Because $R(S) \subseteq S$ and $A \subseteq B \subseteq F(A) \subseteq F(B)$, $F(R(S)) \subseteq F(S)$. Therefore $R(S) \subseteq F(R(S)) \subseteq F(S)$.