

Topologically Clean Distance Fields

Attila G. Gyulassy, Mark A. Duchaineau, *Member, IEEE*, Vijay Natarajan, *Member, IEEE*, Valerio Pascucci, *Member, IEEE*, Eduardo M. Bringa, Andrew Higginbotham, and Bernd Hamann, *Member, IEEE*

Abstract—

Analysis of the results obtained from material simulations is important in the physical sciences. Our research was motivated by the need to investigate the properties of a simulated porous solid as it is hit by a projectile. This paper describes two techniques for the generation of distance fields containing a minimal number of topological features, and we use them to identify features of the material. We focus on distance fields defined on a volumetric domain considering the distance to a given surface embedded within the domain. Topological features of the field are characterized by its critical points. Our first method begins with a distance field that is computed using a standard approach, and simplifies this field using ideas from Morse theory. We present a procedure for identifying and extracting a feature set through analysis of the MS complex, and apply it to find the invariants in the clean distance field. Our second method proceeds by advancing a front, beginning at the surface, and locally controlling the creation of new critical points. We demonstrate the value of topologically clean distance fields for the analysis of filament structures in porous solids. Our methods produce a curved skeleton representation of the filaments that helps material scientists to perform a detailed qualitative and quantitative analysis of pores, and hence infer important material properties. Furthermore, we provide a set of criteria for finding the “difference” between two skeletal structures, and use this to examine how the structure of the porous solid changes over several timesteps in the simulation of the particle impact.

Index Terms—Morse theory, Morse-Smale complex, distance field, topological simplification, wavefront, critical point, porous solid, material science.

1 INTRODUCTION

There exists substantial interest in simulations of particle impact. At the nm scale, particles made of few-thousands of atoms can be used to smooth or create nano-scale relief on a variety of surfaces [17]. At the macroscopic scale, particles of mm size are often used to mimic impact of a variety of projectiles, from bullets to meteoroids [21]. Recently, the Stardust mission explored the craters left by micrometeoroids, reaching sub-micron sizes, in the frame of their comet dust catcher [16]. This dust catcher was made of aerogel, a silica foam where the comet particles, moving at ~ 6 km/s, were slowly decelerated and stored for recovery after the aerogel returned, and were subsequently analyzed at several laboratories worldwide.

Understanding of such deceleration and storage is based on experiments and continuum-scale models of particle impact. Such models require equation of state and other thermodynamic input, and might fail at scales where atomistic effects would preclude continuum coarse graining.

To understand possible limitations of continuum models at the

nanometer scale, we have carried out classical molecular dynamics simulations of the impact of a dense grain into low density foam (25% of solid density). Both the grain and the foam are described by Mishin et al. [20] an embedded atom potential for copper, which reproduces well shock data [3]. The size of the simulated box was 70 nm x 70 nm x 80 nm, and the impact velocity of the 4 nm radius, porosity-free, spherical particle was 5 km/s, which makes the impact supersonic (considering the velocity of sound in copper is 4 km/s). Due to this impact, a crater is formed. The analysis of this data must answer the following questions: How can we quantify the loss of porosity of the material? How does the filament density profile of the material change? What is the portion of the material that is affected by the impact crater? How does the structure around the impact crater change? We answer these questions using two methods that compute a curved skeleton and a clean distance field representation of the data. These methods naturally capture the structure of the material, and hence are ideal for this kind of analysis.

We focus on analysis of the filament structure of a porous solid. The porous solid is represented as a distance field over a volumetric grid that is generated based on distance to the interface surface between solid material and empty space. The interface forms a surface that represents the filaments of the porous solid. The tools we have developed are used to recover the latent structure from the experiment. However, several factors make this difficult. First, the scale of the structure we wish to identify is not known. While scientists may have a general idea of the size and distribution of the structure, any analysis based on such guesses would be skewed. In addition, the analysis tools should be valid for any size and distribution of the structure. Second, there is inherent instability in any choice of criteria for identification of such a structure. There are several classes of methods for identifying curved skeletal structures [8], however, they assume a surface representation is given, and construct a skeleton based on that surface. The interface surface of the distance field representing the porous material ideally can be extracted as surface for isovalue zero. However, any choice of isovalue to select a “base surface” is unstable; Figure 1 shows how small changes in isovalue can produce profound differences between the resulting structures being the basis for analysis. Furthermore, the distance field itself has noise, artifacts from computation, and artifacts from discretizing a continuous function onto a grid. As a result, straightforward analysis of the distance field yields an excess of critical points that do not represent physically meaningful and relevant

-
- Attila G. Gyulassy is with the Institute for Data Analysis and Visualization, Dept. of Computer Science, University of California, Davis, E-mail: aggyulassy@ucdavis.edu.
 - Mark A. Duchaineau is with the Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, E-mail: duchaine@llnl.gov.
 - Vijay Natarajan is with the Dept. of Computer Science and Automation, Supercomputer Education and Research Centre, Indian Institute of Science, Bangalore, E-mail: vijayn@csa.iisc.ernet.in.
 - Valerio Pascucci is with the Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, E-mail: pascucci1@llnl.gov.
 - Eduardo M. Bringa is with the Material Science and Technology Division, Lawrence Livermore National Laboratory, E-mail: bringa1@llnl.gov
 - Andrew Higginbotham is with the Department of Physics, Clarendon Laboratory, University of Oxford, Oxford, E-mail: a.higginbotham1@physics.ox.ac.uk.
 - Bernd Hamann is with the Institute for Data Analysis and Visualization, Dept. of Computer Science, University of California, Davis, E-mail: hamann@cs.ucdavis.edu.

Manuscript received 31 March 2007; accepted 1 August 2007; posted online 27 October 2007.

For information on obtaining reprints of this article, please send e-mail to: tcvg@computer.org.

features of the function.

We use a modified distance field for our analysis that contains a minimal set of critical points corresponding to the features of the function. The *clean distance field* is a distance field constructed with a distance metric that is similar to the Euclidean metric, except where the function must be constrained to ensure proper critical point behavior. A two-dimensional manifold defined by points ϵ -distance from the skeleton structure of this function is a reference for a family of homeomorphisms for all contours the field. Our first method computes the Morse-Smale complex using a modified version of the algorithm presented in [15], and filters the arcs to extract this skeleton structure. This method allows us to choose between levels of analysis, and also hints at the quality of the choice, *i.e.*, we introduce filter operations to allow the scientist to control what is viewed as noise or artifact, and what is viewed as a feature. Our second method generates a clean distance field using a propagation-based approach. We show that our two methods compute the same skeleton, and produce an analysis structure for the porous solid that is stable.

1.1 Related work

In distance fields, critical points correspond to changes in the behavior of isosurface components. For example, for bivariate functions, upon increasing the function value, minima create new isosurface components, maxima destroy components, and saddle points merge or split isosurface components. A purely geometric approach to simplification is able to remove small topological features but does not provide the desired level of control. Considerable work has been done on topological simplification of scalar functions. Initial work focused on simplifying topological features or preserving them while simplifying mesh geometry [6, 12]. Two data structures are commonly used for explicitly storing topological features: *Reeb graphs* and *Morse-Smale (MS) complexes*.

The Reeb graph [22] traces components of contours/isosurfaces as one sweeps through the allowed range of isovalues. In the case of simply connected domains, the Reeb graph has no cycles and is called a *contour tree*. Reeb graphs, contour trees, and their variants have been used successfully to guide the removal of topological features [7, 4, 13, 27, 28, 30]. The MS complex decomposes the domain of a function into regions having uniform gradient flow behavior [24]. It has been used recently to perform controlled simplification of topological features in functions defined on bivariate domains [2, 11], in trivariate domains [14, 15], and for purposes of shape analysis [5]. The MS complex allows the simplification to utilize a global view of the function and its spatial distribution for detecting, ordering and removing features along with the ability to restrict the simplification to a local neighborhood of the non-significant feature. Reeb graph-based simplification methods do not enjoy these benefits. Furthermore, when applied to trivariate functions, they are limited to detecting and simplifying features that are associated with the creation and destruction of isosurface components. These features are represented by pairs of critical points consisting of one saddle and one extremum. The MS complex is able to detect genus changes within the isosurface, which are represented by saddle-saddle pairs. This is crucial in analysis of the porous solid, where features are defined by filaments and tunnels in the isosurface. We use this more comprehensive approach for simplifying scalar functions in three variables.

Sethian computed the distance field from a surface using advancing fronts and a priority queue [23]. These distance field calculations are a numeric solution to the Eikonal equations [29], and their efficiency was improved to linear time in [31]. Simplifying distance fields has also been studied, mostly in the context of simplifying particular isosurfaces. Modification of the scalar field to remove isosurface components was presented in [25], where regions are carved permitting only a fixed number of topology changes to the isosurface. This was refined with a bounded error in [26]. The drawback of this approach is that it allows only a single level of resolution of simplification, and it is not guaranteed to remove all low persistence features. Simplifying a single surface to remove handles was also studied in [13, 30].

Computation of the Morse-Smale complex for volumetric domains




Fig. 1. Dependency of extracted core structure and isovalue. The dark lines indicate the structure that we view as a valid approximation of the actual core structure (curved skeleton) for a specific isovalue. A small change in isovalue can have a dramatic change on the core structure.

was first described in [10]. Efficient algorithms for computation and simplification of 3D Morse-Smale complexes were described in [14], and extended to larger datasets in [15]. We extend these results adding extended controls and analysis tools to focus exploration according to the particular needs of the application.

1.2 Results

We present two methods for the generation of clean distance fields, and we demonstrate their usefulness by finding the filament structure for a porous solid. Our first method uses the Morse-Smale complex to display a distribution of critical point pairs of a standard distance field, and introduces filtering operations to extract any apparent features. We show how the arcs representing the filament structure for the porous solid can be recognized using this analysis tool and how to perform topology-based simplification to find them. Our second method proceeds by propagating an advancing front for the particular task of creating a clean distance field. This is an efficient method that modifies the function itself, and also finds the filament structure for the porous solid. We compare the two techniques in terms of level of control and performance, and we define a set of criteria for determining a meaningful distance measure between the curved skeletons they produce.

Application of these methods to the pore impact dataset reveal an important fact: there is significant densification of the foam below the crater wall, while the structure of the foam outside the immediate crater is unaffected. By constructing the analysis structure for relevant timesteps, and applying our structural comparison technique, we are able to identify the changes in the core structure in a quantitative as well as visual manner.

2 TOPOLOGICAL SIMPLIFICATION

A typical distance field has noise or artifacts from construction, or artifacts from quantization. Critical point analysis on such a function relies on topological simplification, *i.e.*, the ability to identify which critical points represent actual features, and selectively remove those that do not. We use the foundation of Morse theory to achieve this simplification.

A smooth scalar function $f : \mathbb{M} \rightarrow \mathbb{R}$ defined on a smooth three-dimensional manifold \mathbb{M} is a *Morse* function if none of its critical points are degenerate *i.e.*, the Hessian matrix at all critical points is non-singular. The distance field is computed and available to us as a sample over a hexahedral grid. We simulate a perturbation [9, Section 1.4] to ensure that all critical points are non-degenerate and hence identify the given distance field as a Morse function. We use ideas from Morse theory to control explicitly the topology of the distance field. We use the phrase “topology of a scalar function f ” to refer to the topological structure of isosurfaces of f . Critical points of f determine topology changes in isosurfaces as they sweep the domain.

To gain intuition, we first describe the topological simplification as applied to a univariate g , see Figure 2. Critical points (maxima and minima) of g partition the domain into monotonic regions. This partition is stored as a graph whose nodes are the critical points of g and edges represent the monotonic curves. Pairs of critical points identify topological features of the function. The size of each feature is defined as the absolute difference in function value between the two critical points and is called the *persistence* of the critical pair. The smaller-sized features are not significant, probably due to noise in the data,




Fig. 2. Multi-scale analysis of a univariate function. (a) Visualization of the function. (b) Critical points of the function partition the domain into monotonic regions. Pairs of critical points identify features, whose sizes are equal to the difference in function value of the critical points. (c) Small-sized monotonic regions are explicitly identified and removed, leaving behind the “significant” features.

and can be removed explicitly to obtain a global view of the function. Removal of critical pairs can be implemented in a purely combinatorial fashion by updating the graph representation of the partition. A formal Morse theory-based approach to this multi-scale analysis helps in extending the topological feature-simplifying operations to bivariate and trivariate functions.

2.1 Morse-Smale complex

Morse theory studies the relationship between critical points of a Morse function and the topological structure of its domain space [19]. The *Morse Lemma* states that in the neighborhood of a critical point p of f , the function can be rewritten as a quadric

$$f(x, y, z) = f(p) \pm x^2 \pm y^2 \pm z^2.$$

The *index* of p is equal to the number of negative signs in the above expression. Critical points of index 0, 1, 2, and 3 are called *minimum*, *1-saddle*, *2-saddle*, and *maximum*, respectively. This characterization of critical points was transported to piecewise-linear functions by Banchoff [1] and later used by Edelsbrunner et al. [10, 11] to obtain combinatorial algorithms for characterizing critical points and to compute the *Morse-Smale complex* of Morse functions. The Morse-Smale complex (MS complex) decomposes the domain \mathbb{M} into monotonic regions and represents the topological structure of f .

An integral line of f is a maximal path in \mathbb{M} whose tangent vectors are equal to the gradient of f at every point of the path. Each integral line has an origin and a destination at critical points of f where the gradient becomes zero. A cell in the MS complex is a set of all integral lines that share a common origin and destination. For example, the 3-dimensional cells of the MS complex are sets of integral lines that originate at a given minimum and terminate at an associated maximum. The cells of dimensions 3, 2, 1, and 0 of the MS complex are respectively called *crystals*, *quads*, *arcs*, and *nodes*.

2.2 Multi-scale analysis

The Canceling Handles Theorem [19, Section 3.4] leads to an algorithm for simplifying the MS complex and hence the topology of f . This theorem essentially states the following:

CANCELING HANDLES THEOREM. Critical points can be destroyed in pairs that differ in index by one and are connected by an arc in the MS complex. The cancellation is numerically realized by a local perturbation of the gradient field.

Given an ordered list of critical pairs, the MS complex can be simplified by canceling the critical point pairs in succession using a combinatorial algorithm developed by Gyulassy et al. [14]. This algorithm constructs an artificial complex by introducing “dummy” critical points at the vertices of a barycentric subdivision of the input cube mesh. The index of criticality of the dummy critical point is equal to the dimension of the mesh element and the function value at the dummy nodes are chosen such that they have an infinitesimally small persistence (persistence being the absolute difference in function value of the cancelled pair or critical points). The MS complex is obtained in a pre-processing step by canceling all ϵ -persistence critical pairs (ϵ is infinitesimally small). The critical pairs are ordered based on their persistence and given a persistence threshold, all critical pairs below this

threshold are canceled. The critical pairs are classified into two categories: *saddle-extremum* and *saddle-saddle*. Saddle-extremum pairs consist of a minimum and a 1-saddle or a maximum and a 2-saddle. The two types of saddle-extremum pairs are dual to each other. The duality is given by a negation operator acting on the function that maps critical points of index i to index $3 - i$. A saddle-saddle pair consists of a 1-saddle and a 2-saddle.

2.3 Filter-driven simplification

Meaningful and important features of a given function are not always captured by the notion of persistence. For example, a scientist may be interested in the function behavior within a region enclosed by certain isosurfaces. In this case, simplification should ideally preserve the topological structure of the isosurface components while removing noise in the volumetric region inside and outside. Extrema with function value within a given range may correspond to relevant features, and in this case simplification should leave these extrema unaffected. Both cases arise for the distance fields that we study. In fact, features may arise in locations not initially predicted. The MS complex is a useful tool in identifying such features since it provides a full characterization of the gradient flow behavior (when viewing the function’s gradient as a flow field). Therefore, analysis of the critical point pairs and arcs of the complex can lead to better understanding of the actual locations of the features, and where to apply topological simplification.

The porous solid dataset suffers from the fact that the distance field was created from an interface surface that is unstable. A small change in the selection of this surface could lead to a profound difference in the topology. However, by having relaxed notions of the exact location of this interface, we can overcome the instability and produce a result that is invariant under small changes in selection of the interface surface.

We use several filters to direct our simplification process to preserve relevant features in the data (relevance understood here as user-specified features for a particular application). A filter specifies the arcs of the MS complex that are to be removed from the list of candidates for cancellation. Any filtering requirements can be met with the following three conditions:

- i* Arcs that have their lower, upper, both, or neither end points in a given range of function values.
- ii* Arcs that cross a given isovalue.
- iii* Arcs whose lengths lie within a given range.

These criteria, or combinations thereof, designate a wide range of features.

We show an example where the distribution of critical point pairs help distinguish between actual features and artifacts in Figure 3. In this example, we created a distance field using a standard approach as the signed distance from the shells of a set of atoms distributed along a spiral and a sinusoidal curve. The atoms are placed at random along these curves, and additional “noise” atoms are added throughout the data. We compute the Morse-Smale complex for this distance field. We wish to extract the curved skeleton from this distance field, without knowing *a priori* the details about how the distance field was constructed. Intuitively, we can guess that the features will be represented by 2-saddle - maximum pairs where the maximum will have large




Fig. 3. The isosurface for isovalue zero of the initial distance field (a). We compute the Morse-Smale complex of this field (b), and apply filtering to extract the stable core structure (c).

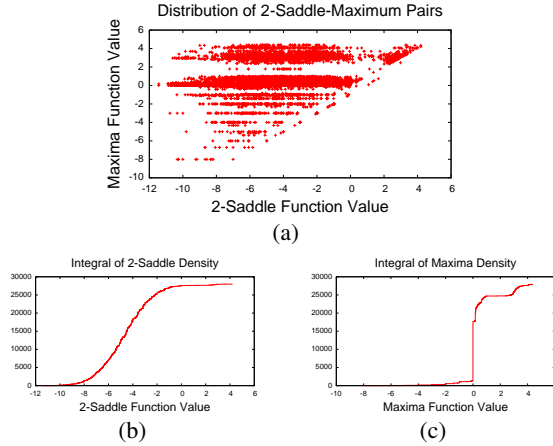


Fig. 4. The distribution of 2-saddle - maximum pairs (a). Each pair is plotted as a point whose coordinates are the function values of the 2-saddle and the maximum. We integrate along the x-axis (b) and y-axis (c) to see the density distribution of 2-saddles and maxima.

positive value and the 2-saddle has large negative value. The distribution of critical point pairs, illustrated in Figure 4, suggests certain stable thresholds for “important” maxima and 2-saddles. In particular, these pairs are those in the upper-left corner of the scatter plot. Flat regions in the integral along each axis reveal that the stable threshold for maxima is two, and the stable threshold for 2-saddles is -1, where this curve first starts to flatten out. By cancelling all arcs that do not entirely cross the range $[-1, 2]$, we remove the artifacts and noise. The core structure is extracted as the 2-saddle - maximum arcs that remain and are entirely contained in the isosurface for isovalue zero.

2.4 Fingers

Persistence-based simplification of the MS complex can result in configurations that cannot be reduced. Such a situation arises when a cell of the complex contains three pairs of critical pairs, none of whom can be canceled because of an obstruction, as shown in Figure 5(a). Cancellation of any of the three critical pairs results in an invalid MS complex as shown by Gyulassy et al. [14]. Canceling the saddle-extremum pair leads to a *strangulation* of the cell whereas canceling the saddle-saddle pair results in the creation of a *pouch*. The cell cannot be removed by any sequence of cancellations. These configurations, called *fingers*, are artifacts that result from our choice of order of cancellations in flat regions of the function. In our construction, we introduce flat regions throughout the function to create the evenly spaced artificial complex. Therefore, these fingers can accumulate in large numbers as shown in Figure 5(b). In fact, the maximum persistence within such a finger may be much smaller than the persistence filter and yet the small feature cannot be removed. While the structure of the complex remains combinatorially sound, the fingers add to visual clutter, and should be removed. This is especially important in analysis of the porous solid, since fingers might erroneously indicate structural components of the distance field.




Fig. 5. Certain orders of cancellation result in finger-like configurations. (a) A single finger contains three critical pairs none of which can be legally canceled even if they have low persistence. (b) Fingers could accumulate leading to visual artifacts.




Fig. 6. The geometric location of an integral line in a flat region is arbitrary. Therefore we shortcut the path, and find the shortest path through the flat region.




Fig. 7. Artifacts from the construction of the complex include flat regions inside every voxel. Therefore, an integral line can take any valid path from the point of entry in a voxel to the point of exit. We again remove kinks in the integral line by shortcutting the path.

We use the sliding window algorithm for construction of the MS complex presented in [15] with slight modifications. We prevent the creation of fingers by reordering the cancellations in the preprocessing stage of this algorithm. All saddle-extremum cancellations are scheduled before the saddle-saddle cancellations to ensure that none of the ϵ -persistent saddle-extremum pairs remain. While obstructions can still develop, they are resolved by future saddle-saddle cancellations. Additionally, we perform all possible cancellations within a single slice of the data before canceling critical pairs that span multiple slices. In particular, this removes all minimum-1-saddle pairs within the slice. Since 2-saddle-maximum pairs span multiple slices, this reordering prevents the simultaneous existence of un-cancellable minimum-1-saddle and 2-saddle-maximum pairs.

2.5 Arc smoothing

The procedure described in [14] for constructing the MS complex from the artificial complex results in arcs that contain geometric artifacts: the sequence of line segments representing the arc may differ significantly from the location of the corresponding integral line. This discrepancy occurs because the artificial complex introduces several small regions of constant value, “flat regions,” to the function. Integral lines are not uniquely defined in these flat regions, and therefore the arcs we produce may wander before resuming the path of steepest ascent, see Figure 6. Figure 7 shows how these flat regions can introduce sharp spikes in the arcs. In fact, while the combinatorial structure of the complex is correct, the geometry of the arcs may be off from the integral lines by one voxel in any direction. For use in the analysis of the distance field, we want the arcs to behave like simple curves that can be represented as a sequence of line segments.

The spikes we introduce are the result of cancellations in flat regions. Previous approaches [2] perturb the function to avoid flat regions. However, the necessity to reorder cancellations to avoid creation of fingers in our approach prohibits such an approach. Instead, we shift the arcs towards the integral line using shortcuts. The spikes in the arcs have a unique property that they occur entirely within a single hexahedral cell, being a unit cube (or voxel) in our case, around




Fig. 8. Legal and illegal 3x3x3 neighborhood configurations: (a) 2x2 edge legal and illegal cases, showing only one of four rotations; (b) vertex neighborhoods need only test for diagonal cases; (c) tubes should not be created or destroyed.

the arc. These unit cubes have the property that they have unique points where the arc enters p_{entry} and where the arc exits p_{exit} . We consider the slope of any line with end points x_0 and x_1 in the cube as $(f(x_1) - f(x_0)) / \|x_1 - x_0\|$ and therefore, given p_{entry} and p_{exit} , the maximum slope is in the direction that minimizes the distance. Hence we use a shortest-path algorithm within unit cubes to shortcut the spikes and yield a smoother arc.

3 CLEAN DISTANCE FIELD PROPAGATION ALGORITHM

The second method of constructing topologically clean fields uses a variant on queue-based signed distance field construction for regular volume grids, modified to allow only manifold and topologically clean isocontour propagations. The original signed distance propagation algorithm is from Laney et al. [18]. In the new approach, only distances are kept in the construction state, not actual closest points. Another change is that a fast priority queue is used, based on distance, in place of a first-in-first-out queue.

The state of the distance propagation includes a 3D regular grid of distances per cell, and a priority queue using fast bucket sort of the potentially updating cells on the propagation front. The algorithm works in two passes, one for the positive distances and a second symmetric pass for the negative distances. The algorithm takes as input a conventional signed distance field that has already been generated for reference.

For the positive clean distance propagation, all negative distance entries from the conventional signed distance field are left as is, while the positive values are initially replaced with infinity. All positive infinity cells are checked for legal update capability. If the cell could be legally updated, the cell is put on the priority queue. One by one, the smallest distance cell on the priority queue is dequeued and updated (its distance is stored in the clean distance field), and all its 26 neighbor cells are tested for legal update capability. Any neighbor that is either already on the update queue or could be legally updated is queued with a new update distance as its priority. This processing proceeds until the priority queue is empty. It remains to describe the test for legality of updates and the update distance values.

The legality of update, and the distance value upon update, are determined by the following procedure for cell i . First, the maximum of all non-infinity values of the 26 neighbors, f_{max} , is determined. Next the minimum neighbor j_{min} with value f_{min} is determined. The smallest potentially legal update distance value is then $f_{up} = \max\{f_{min} + \text{dist}(i, j_{min}), f_{max} + \epsilon\}$, where $\text{dist}(i, j_{min})$ is the distance from the center of cell i to the center of minimum neighbor cell j_{min} , and where ϵ is an infinitesimal positive value (in practice, the step to the next IEEE floating point value). This minimum potential update value f_{up} is then tested for legality.

Legality of propagation is based on a cell-face model of contours. In other words, the field is thought of as piecewise constant with distinct values per cell, with contours formed from sets of cell faces. The legality test for an update uses a bit mask for the 3x3x3 neighborhood around the potentially updating cell, where each of the 27 bits is set to one if the corresponding cell is infinity, and zero otherwise. This bit mask is then tested to ensure that manifold propagation of the contour faces occurs, and that the contour topology is unchanged. This can be thought of as incrementally adding cubes to the “in” set of a solid such that the surface of the solid stays manifold and keeps the same number

of handles and components.

This incremental guarantee is obtained by ensuring that: all twelve edges of the added cube remain manifold with zero or two faces incident, see Figure 8(a); that all eight vertices do not have a diagonal configuration see Figure 8(b); that a single box or void is not created; that a homogeneous (all in or out) region is not created; and that an axial tube is not created or destroyed, see Figure 8(c). Other than these restrictions on legal updates, the distances propagate as in a conventional distance field update, with forced strictly increasing values for cells relative to their neighbors. This implies that cells that are delayed for inclusion due to legality constraints will have a higher magnitude distance value than if they were introduced at the traditional (non-clean) step in the queue processing.

The testing for legality can be implemented through 3x3x3 neighborhood bit mask manipulations with small table lookups per update cell edge (four bit lookup index), vertex (eight bit lookup index) and tube end configuration (six bit lookup index). This is much faster to evaluate than a large sequence of individual assertions, and requires far less memory (and is thus far more cache friendly) than using a giant table with 27-bit lookup indices for the full 3x3x3 neighborhood state.

4 RESULTS

We demonstrate the usefulness of our two procedures for the generation of clean distance fields by finding the filament structure for a porous solid. Each method produces an output representing the curved skeleton structure of the clean distance field, which is the core structure of the porous solid. We compare the two methods by defining criteria to determine the similarity of the resulting core structures. The porous solid dataset is a standard distance field derived as the signed Euclidean distance from the shells of atoms in the simulation.

4.1 Core Structure of a Porous Solid

We compute the core structure of a porous solid using both methods. First, we compute the Morse-Smale complex using the incremental algorithm presented in [15], with the finger removing and arc smoothing modifications. Then, similar to our previous example shown in Figure 3, we analyze the critical point pairs, and filter the arcs to extract the core structure. The full complex shown in Figure 9, is used to plot the distribution of 2-saddle - maximum pairs, shown in Figure 10. The features we are interested in are the arcs that connect a low 2-saddle to a high maximum; these critical point pairs are in the top left of the distribution. Flat regions in the integral along each axis reveal that the stable threshold for maxima is 1.5, and the stable threshold for 2-saddles is -0.8 . By cancelling all arcs that do not entirely cross the range $[-0.8, 1.5]$, we remove the artifacts and noise. The core structure is extracted as the 2-saddle - maximum arcs that remain after simplification and are entirely contained in the isosurface for isovalue zero. For this particular application, we are interested in the connectivity of the porous solid, therefore we omit arcs that are connected to the structure at only one endpoint from the final core structure.

The second method creates a clean distance field starting with an isosurface at a chosen threshold value, shown in Figure 11. We use -0.8 as the stable threshold found through analysis of the critical point pairs. Due to the propagation of the topology-preserving front used in this method, all “dangling” arcs are retracted, leaving the same structure we found using the first method.

The results were generated using an off-the-shelf personal computer, a 3.4GHz Pentium 4, with 2 Gb of memory. The porous solid was represented as real-valued samples on a $230 \times 230 \times 375$ regular grid. The total time required for computation of the initial MS complex was 6 hours 32 minutes and 45 seconds. Additional processing to attain the graph structure took 32 seconds. The total time required to create the clean distance field using the propagation method starting with the same input dataset was 91 seconds. After the initial computation, exploration and further simplification can be done interactively. For large datasets, computation of the full MS complex is not possible; however, we can still perform analysis on a subset of the data to attain the distribution of critical point pairs, assuming features are distributed




Fig. 9. The initial computation of the MS complex for the full dataset (a) is simplified revealing the graph structure (b) of the porous solid.

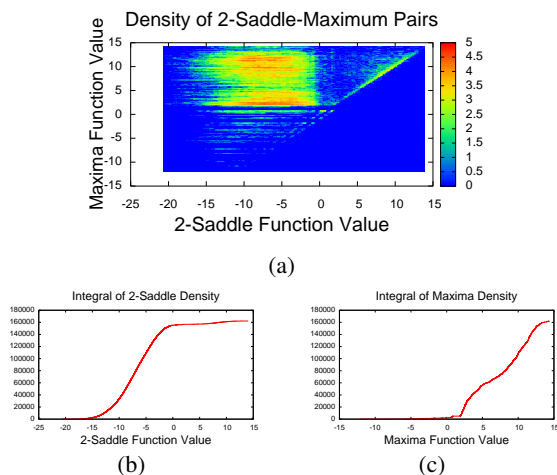


Fig. 10. The distribution of the 2-Saddle-Maximum pairs (a). The color scale shows the number of occurrences of 2-Saddle-Maximum pairs where the 2-Saddle has function value along the x -axis, and the Maxima has function value along the y -axis. We integrate along the x -axis (b) and along the y -axis (c) to find stable threshold values.

nearly uniformly in the dataset. The clean distance field method is an efficient implementation that can be applied to large data, and benefits from such an analysis.

4.2 Comparison

We present a qualitative comparison of the two results using a visualization overlaying the two core structures in Figure 12. The “struts” connecting the two overlaid graphs represents the connection to the closest points of each with respect to the other.

We also present a set of criteria as a quantitative measure of difference between the structures. Specifically, we have used these criteria:

- i* Hausdorff distance - the maximum geometric distance between points on one graph and their closest neighbor on the other. To obtain a symmetric measure, the geometric distance is computed in both directions and the larger value chosen.
- ii* Average distance between closest pairs on the two graphs.
- iii* Number of simple cycles in each graph, to estimate connectivity.
- iv* Total length of edges in each graph.

A number of differences exist in the results produced by each method. Differences in the geometry result from the fact that the edges from the MS complex are restricted to edges of the grid, while edges from the clean distance method are smoothed in the direction of the gradient. This factor, however, contributes only a small fraction of the difference between the two methods. A number of cases arise where the connectivity of the two graph structures is different, however, this




Fig. 11. The initial isosurfaces (a) reveal noise. Computing the clean distance field removes small isosurface components and reveals the filament structure (b).




Fig. 12. Comparative visualization of the results obtained by the two methods. The teal structure indicates points on the graph returned by the method based on the MS complex, and the yellow structure indicate points on the graph returned by the propagation method. The red lines indicate the closest point from each point on the first graph to the second graph, and the blue lines show the closest point from each point on the second graph to the first graph.

is a result of properties of each algorithm, and not due to instability. In particular, the largest difference between the two algorithms was retraction of “dangling” arcs. A small loop at the end of such an arc would prevent retraction of that arc. These small differences occur due to the necessity of selecting a single original isosurface as the start of the propagation method. Overall, these differences contribute less than one percent of the length of the core structures.

Table 1 shows the results of comparing the two methods for the porous structure dataset.

Metric	MS-c method	Prop method
Hausdorff	31.5	33.53
average distance	1.7	1.4
number of cycles	372	304
total length	22239	19002

Table 1. Comparison results for the two procedures.

4.3 Time-Dependent Impact Data

We have used our methods to explore a simulated dataset of a particle impact on the porous material at several timesteps. By computing the clean distance field, we can obtain the density of the porous solid as the ratio of the number of sample points interior to the zero isosurface to the number of sample points outside. The clean distance field ensures that all sample points identified in this way contribute to the structure of the porous material. The structural analysis and comparison between time steps allow us to obtain an important result: there is significant densification of the foam below the crater wall. Such analysis provides a simple, quantitative answer to the amount of densification, as shown in Figure 14 (a,b). In this figure we show the density profile at different times as a function of depth for slices of (a) the whole sample, and (b) the center of the sample, including the crater. It can be seen in (b) that the density increases by a factor of two

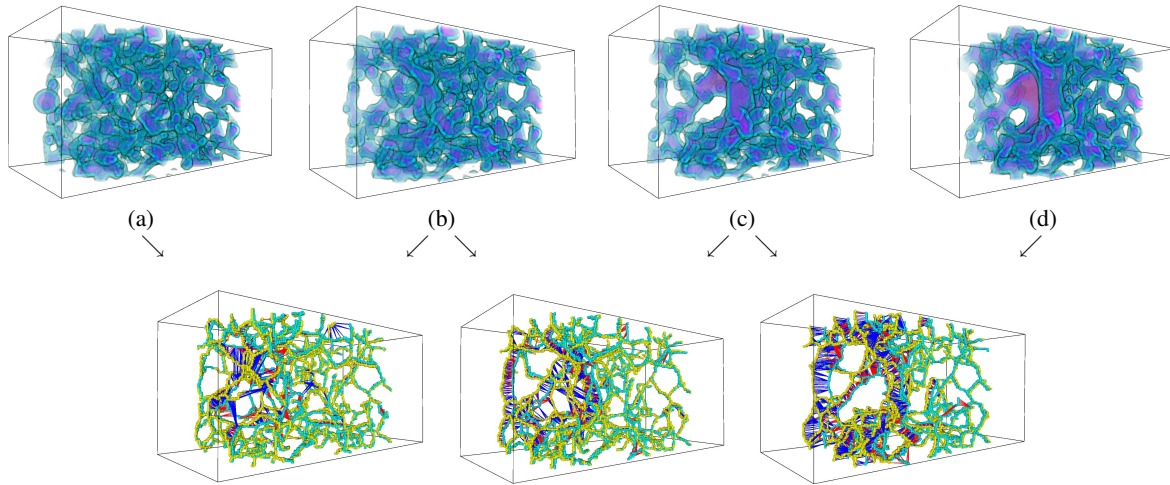


Fig. 13. *Top*: A volume rendering of the impact of the ball entering the porous solid from the left at time step 500 (a), time step 12750 (b), time step 25500 (c), and time step 51000 (d). *Bottom*: We compare the core structures of consecutive time steps. The yellow dots represent the core structure of the initial time step, and the teal dots represent the core structure of the next time. The closest arcs between in the core structures at the different time steps are connected via blue and red line segments. The length of these segments corresponds to the displacement of the arc.




Fig. 14. We compute the filament density of the material as the ratio of samples inside the interface surface to those outside the interface surface at each time step. We compute this density profile for the entire data (a), and a cylinder through the impact crater (b).

below the crater.

We also have computed the core structure using the MS complex for each timestep. Using the distance measures from Section 4.1, we can compare qualitatively how the structure of the material changes between timesteps. Figure 13 is a visual depiction of the displacement of filament segments at different times after impact. Note the large displacements near the crater, and the nearly zero displacement well below the crater.

In both figures it can be seen that density barely changes in the bottom one-third of our sample. This is partly due to the fact that the foam is extremely efficient at absorbing the impact shock wave. Key statistics of the core structure for each time steps are summarized in Table 2.

Metric	t=500	t=12750	t=25500	t=51000
# Cycles	762	340	372	256
Total Length	34756	24316	23798	18912

Table 2. Statistics for each timestep.

The ratio of cycle counts before and after the impact supports this observation, as approximately two-thirds of the cycles are destroyed. The ratio of the total length of the filaments before and after the particle impact implies that volume of material displaced by crater is approximately one-half the volume of the rest of the material. Since this ratio is fairly close to the ratio of the cycle counts, we can say that the majority of the filaments that were broken happened to be in the interior of the crater. The sum of the Hausdorff distances between the timesteps is 98.6, giving the maximum distance that any element of the

material travelled during the impact. This number is surprisingly high, corresponding to the entire depth of the crater; it indicates that the material of the filaments first hit by the particle was displaced along the trajectory of the particle. The average distance between closest pairs in the graphs of the consecutive timesteps was less than 5.0, indicating that the displacement did not propagate into the material, outside the direct path of the particle.

Densification of the foam will vary as a function of impact velocity, and a quantitative characterization of such a function might help to narrow down such velocities when they are not known. In addition, the fact that the foam is getting denser would change, for large particle fluxes, the foam’s mechanical and thermal properties.

5 CONCLUSIONS

We have presented two methods for the construction of topologically clean distance fields. We used these methods to extract, characterize, and visualize relevant filament structures in a porous material. The analysis of critical point pairs of the MS-complex eliminates most of the uncertainty and instability associated with traditional methods, allowing the identification of a consistent and stable core structure through filtering operations. The propagation method is an efficient and scalable method that modifies the function itself to create a clean distance field with the same topology. We showed that the our two methods produce similar core structures for the porous solid. The ability to extract a stable and consistent core structure allowed us to make comparisons across time steps for the particle impact data, and extract meaningful results about the material properties.

ACKNOWLEDGEMENTS

Thanks to E. Taylor for help with the cratering molecular dynamics simulations. Attila Gylassy has been supported by a Student Employee Graduate Research Fellowship (SEGRF), Lawrence Livermore National Laboratory. This work was supported by the National Science Foundation under contract ACI 9624034 (CAREER Award), through the Large Scientific and Software Data Set Visualization (LSSDSV) program under contract ACI 9982251, and a large Information Technology Research (ITR) grant. This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.

REFERENCES

- [1] T. F. Banchoff. Critical points and curvature for embedded polyhedral surfaces. *Am. Math. Monthly*, 77:475–485, 1970.

- [2] P.-T. Bremer, H. Edelsbrunner, B. Hamann, and V. Pascucci. A topological hierarchy for functions on triangulated surfaces. *IEEE Transactions on Visualization and Computer Graphics*, 10(4):385–396, 2004.
- [3] E. M. Bringa, J. U. Cazamias, P. Erhart, J. Stolken, N. Tanushev, B. D. Wirth, R. E. Rudd, and M. J. Caturia. Atomistic shock hugoniot simulation of single-crystal copper. *Journal of Applied Physics*, 96(7):3793–3799, 2004.
- [4] H. Carr, J. Snoeyink, and M. van de Panne. Simplifying flexible isosurfaces using local geometric measures. In *Proc. IEEE Conf. Visualization*, pages 497–504, 2004.
- [5] F. Cazals, F. Chazal, and T. Lewiner. Molecular shape analysis based upon the morse-smale complex and the connolly function. In *SCG '03: Proceedings of the nineteenth annual symposium on Computational geometry*, pages 351–360, New York, NY, USA, 2003. ACM Press.
- [6] Y. Chiang and X. Lu. Progressive simplification of tetrahedral meshes preserving all isosurface topologies. *Computer Graphics Forum*, 20(3):493–504, 2003.
- [7] P. Cignoni, D. Constanza, C. Montani, C. Rocchini, and R. Scopigno. Simplification of tetrahedral meshes with accurate error evaluation. In *Proc. IEEE Conf. Visualization*, pages 85–92, 2000.
- [8] N. D. Cornea, D. Silver, and P. Min. Curve-skeleton applications. *IEEE Conf. Visualization*, 0:95–102, 2005.
- [9] H. Edelsbrunner. *Geometry and Topology for Mesh Generation*. Cambridge Univ. Press, England, 2001.
- [10] H. Edelsbrunner, J. Harer, V. Natarajan, and V. Pascucci. Morse-Smale complexes for piecewise linear 3-manifolds. In *Proc. 19th Ann. Sympos. Comput. Geom.*, pages 361–370, 2003.
- [11] H. Edelsbrunner, J. Harer, and A. Zomorodian. Hierarchical Morse-Smale complexes for piecewise linear 2-manifolds. *Discrete and Computational Geometry*, 30(1):87–107, 2003.
- [12] T. Gerstner and R. Pajarola. Topology preserving and controlled topology simplifying multiresolution isosurface extraction. In *Proc. IEEE Conf. Visualization*, pages 259–266, 2000.
- [13] I. Guskov and Z. Wood. Topological noise removal. In *Proc. Graphics Interface*, pages 19–26, 2001.
- [14] A. Gyulassy, V. Natarajan, V. Pascucci, P.-T. Bremer, and B. Hamann. Topology-based simplification for feature extraction from 3d scalar fields. In *Proc. IEEE Conf. Visualization*, pages 535–542, 2005.
- [15] A. Gyulassy, V. Natarajan, V. Pascucci, P. T. Bremer, and B. Hamann. A topological approach to simplification of three-dimensional scalar fields. *IEEE Transactions on Visualization and Computer Graphics (special issue IEEE Visualization 2005)*, pages 474–484, 2006.
- [16] F. Horz, R. Bastien, J. Borg, J. P. Bradley, J. C. Bridges, D. E. Brownlee, M. J. Burchell, M. Chi, M. J. Cintala, Z. R. Dai, Z. Djouadi, G. Dominguez, T. E. Economou, S. A. J. Fairey, C. Floss, I. A. Franchi, G. A. Graham, S. F. Green, P. Heck, P. Hoppe, J. Huth, H. Ishii, A. T. Kearsley, J. Kissel, J. Leitner, H. Leroux, K. Marhas, K. Messenger, C. S. Schwandt, T. H. See, C. Snead, I. Stadermann, Frank J., T. Stephan, R. Stroud, N. Teslich, J. M. Trigo-Rodriguez, A. J. Tuzzolino, D. Troadec, P. Tsou, J. Warren, A. Westphal, P. Wozniakiewicz, I. Wright, and E. Zinner. Impact Features on Stardust: Implications for Comet 81P/Wild 2 Dust. *Science*, 314(5806):1716–1719, 2006.
- [17] Z. Insepov and I. Yamada. Molecular dynamics simulation of cluster ion bombardment of solid surfaces. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 99:248–252, May 1995.
- [18] D. Laney, M. Bertram, M. Duchaineau, and N. Max. Multiresolution distance volumes for progressive surface compression. *Proceedings of 3D Data Processing Visualization and Transmission*, 00:470–479, 2002.
- [19] Y. Matsumoto. *An Introduction to Morse Theory*. Amer. Math. Soc., 2002. Translated from Japanese by K. Hudson and M. Saito.
- [20] Y. Mishin, M. J. Mehl, D. A. Papaconstantopoulos, A. F. Voter, and J. D. Kress. Structural stability and lattice defects in copper: Ab initio, tight-binding, and embedded-atom calculations. *Phys. Rev. B*, 63(22):224106, May 2001.
- [21] L. E. Murr, S. A. Quinones, E. F. T., A. Ayala, O. L. Valerio, F. Hrz, and R. P. Bernhard. The low-velocity-to-hypervelocity penetration transition for impact craters in metal targets. *Materials Science and Engineering*, 256:166–182, November 1998.
- [22] G. Reeb. Sur les points singuliers d’une forme de pfaff complètement intégrable ou d’une fonction numérique. *Comptes Rendus de L’Académie ses Séances, Paris*, 222:847–849, 1946.
- [23] J. A. Sethian. A fast marching level set method for monotonically advancing fronts. *Proc. nat. Acad. Sci*, 93(4):1591–1595, Sept 1996.
- [24] S. Smale. On gradient dynamical systems. *Ann. of Math.*, 74:199–206, 1961.
- [25] A. Szymczak and J. Vanderhyde. Extraction of topologically simple isosurfaces from volume datasets. In *Proc. IEEE Conf. Visualization*, pages 67–74, 2003.
- [26] A. Szymczak and J. Vanderhyde. Simplifying the topology of volume datasets: an opportunistic approach. Technical Report 09, 2005.
- [27] S. Takahashi, G. M. Nielson, Y. Takeshima, and I. Fujishiro. Topological volume skeletonization using adaptive tetrahedralization. In *GMP '04: Proceedings of the Geometric Modeling and Processing 2004*, page 227, Washington, DC, USA, 2004. IEEE Computer Society.
- [28] S. Takahashi, Y. Takeshima, and I. Fujishiro. Topological volume skeletonization and its application to transfer function design. *Graphical Models*, 66(1):24–49, 2004.
- [29] J. Tsitsiklis. Efficient algorithms for globally optimal trajectories. *Automatic Control, IEEE Transactions on*, 40(9):1528–1538, Sept 1995.
- [30] Z. Wood, H. Hoppe, M. Desbrun, and P. Schröder. Removing excess topology from isosurfaces. *ACM Transactions on Graphics*, 23(2):190–208, 2004.
- [31] H. Zhao. A fast sweeping method for eikonal equations. *Mathematics of Computation*, 74:603–627, Sept 2004.