

A Narrow Band Level Set Method for Surface Extraction from Unstructured Point-based Volume Data

Paul Rosenthal

Vladimir Molchanov

Lars Linsen

Jacobs University, Bremen, Germany

{p.rosenthal, v.molchanov, l.linsen}@jacobs-university.de

ABSTRACT

Level-set methods have become a valuable and well-established field of visualization over the last decades. Different implementations addressing different design goals and different data types exist. In particular, level sets can be used to extract isosurfaces from scalar volume data that fulfill certain smoothness criteria. Recently, such an approach has been generalized to operate on unstructured point-based volume data, where data points are not arranged on a regular grid nor are they connected in form of a mesh. Utilizing this new development, one can avoid an interpolation to a regular grid which inevitably introduces interpolation errors. However, the global processing of the level-set function can be slow when dealing with unstructured point-based volume data sets containing several million data points.

We propose an improved level-set approach that performs the process of the level-set function locally. As for isosurface extraction we are only interested in the zero level set, values are only updated in regions close to the zero level set. In each iteration of the level-set process, the zero level set is extracted using direct isosurface extraction from unstructured point-based volume data and a narrow band around the zero level set is constructed. The band consists of two parts: an inner and an outer band. The inner band contains all data points within a small area around the zero level set. These points are updated when executing the level set step. The outer band encloses the inner band providing all those neighbors of the points of the inner band that are necessary to approximate gradients and mean curvature. Neighborhood information is obtained using an efficient k d-tree scheme, gradients and mean curvature are estimated using a four-dimensional least-squares fitting approach.

Comparing ourselves to the global approach, we demonstrate that this local level-set approach for unstructured point-based volume data achieves a significant speed-up of one order of magnitude for data sets in the range of several million data points with equivalent quality and robustness.

Keywords: Level sets, unstructured point-based volume data, isosurface extraction.

1 INTRODUCTION

Many modern technologies and methods generate volume data that is not gridded anymore. The data points can have an arbitrary distribution without any connectivity. A major group of such data stem from numerical simulations of natural phenomena. They are carried out on unstructured point-based data to gain a high flexibility and provide new insights by additional degrees of freedom. Many physical simulations are, for instance, carried out as smoothed particle hydrodynamics simulations with millions of particles. Such simulations allow for the reproduction of complex natural phenomena by not only simulating the evolution of data at the sample points but also simulating the flow of the sample points under respective forces. Hence, data points move over time, change their positions and neighborhoods, and are distributed with a highly varying density. Such a flexible data structure saves computation time when using high sample density only in those regions, where action takes place.

Level-set methods have a large variety of applications and, in particular, have entered the fields of image processing and scientific visualization for extracting features from scalar data. We are mainly interested in the aspect of segmenting scalar volume data. Many algorithms and approaches with different modifications of the main level-set idea exist. Most of the algorithms address a specific problem or a specific type of data. Typically, the algorithms operate on hexahedral cells and a given initial level-set function is modified to explicitly or implicitly minimize a given energy functional.

Recently, we have presented an approach generalizing the basic idea of level sets to unstructured point-based volume data [19]. It operates directly on unstructured data, i. e., we do not resample the data over a structured grid nor are we generating global or local polyhedrizations. By doing so, resampling errors are avoided, which are inevitably introduced when resampling unstructured data over a regular structured grid using scattered data interpolation techniques.

The level-set method involves the evolution of a function using an iterative numerical integration scheme. At each iteration step, the approach typically induces complex calculations and the evaluation of partial differential equations at each sample location and each iteration step. When dealing with large data sets and small

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

time-integration steps, this can lead to enormous computation times until convergence of the level-set process. The calculations for unstructured point-based data are yet more complex than those in the gridded case, slowing down the process even more. To overcome this problem, narrow-band methods have been proposed in recent years. The main idea is to update the level-set function only in a narrow band of interest and neglect all other sample points, saving a major number of computations. Narrow-band methods have turned out to be very useful for fast yet correct level-set computations on gridded data. To our knowledge, no algorithm exists yet that combines the accurate computations of level-sets with the efficiency of a narrow-band localization directly applied to unstructured point-based volume data. We propose such a method for smooth isosurface extraction based on a fast local level-set approach.

We initialize the level-set function at the sample point positions and extract the zero level set. This zero level set is being transformed into the desired surface by application of the iterative level-set process to the level-set function. Therefore, we can restrict the process to a narrow band of sample points around the current zero level set.

The main ideas of our local level-set method and the structure of the whole visualization pipeline are introduced in Section 3. The creation of the narrow band and the update of the level-set function in the narrow band are explained in Sections 4 and 5, respectively.

The actual level-set process applied to the narrow band and the required approximation of gradient and mean curvature for unstructured point-based volume data using a four-dimensional least-squares method is described in Section 6. Finally, results and their discussion, including the analysis of computation times and detailed comparisons, are provided in Section 7.

2 RELATED WORK

The fast visualization of large unstructured point-based volume data sets is a challenging task, even more when dealing with highly varying sample point densities. The generation of a polyhedrization [4] of the unstructured data points can be very slow for large data sets and is thus not practicable. Hence, the most common way of dealing with unstructured point-based volume data is to resample to a structured grid using scattered data interpolation techniques [7]. Subsequently, a large variety of well-known methods like isosurface extraction, region-growing methods, and level-set methods can be applied to gridded data to generate the desired visualizations. Level-set methods, in particular, are almost without exception executed on regular hexahedral grids what facilitates discrete derivative computations.

Unfortunately, such resampling steps to a regular grid always introduce inaccuracies, which heavily depend on the resolution of the grid and the point density of

the data set. In our case, when dealing with data sets with highly varying point density, the interpolation error can be enormous when using regular grids that fit today's commodity hardware memory constraints. Adaptive grids can reduce the error, but the more adaptive it gets the more complicated the processing becomes and we are looking into data sets with a difference of three orders of magnitude in point density. This raises the desire to directly operate on unstructured point-based volume data.

The original idea of level sets is to implicitly represent a surface as the solution of an equation with respect to an underlying scalar field. The (dynamic) level-set methods go back to Sethian and Osher [15, 16, 21], who first described the evolution of such a closed hypersurface by application of PDEs to the underlying scalar field.

Due to their flexibility and capability level-set methods have developed rapidly in the recent years. Many different approaches exist and the range of application areas is wide. Breen et al. [2] presented a general framework for level-set segmentation of a large variety of regular data sets. Museth et al. [14] use a level-set method to segment non-uniform data sets, where non-uniform data denotes unions of multiple regular data sets with non-uniform resolutions. Enright et al. [5] apply a level-set approach to an octree-based adaptive mesh. Many other approaches, e.g. [22], exist that target a huge variety of level-set segmentation tasks on different structured data sets.

In terms of unstructured data, the particle level-set methods [8] use free particles during the level-set computations, but still need an underlying structured grid to compute the motion of the particles. We recently presented the first approach, directly processing level sets on unstructured point-based volume without any grid calculation or reconstruction of the scalar field [19]. We initialize the level-set function only at sample positions and carry out the whole level-set evolution at these locations. One drawback of the method is the relatively low speed when compared to state-of-the-art level-set methods for gridded data.

The largest gain in computation speed for level-set methods was achieved with the introduction of local level-set methods by Adalsteinsson and Sethian [1]. These methods, also called narrow-band methods, have been developed rapidly in the last years [12, 17] and have also been transformed to work on today's fast and parallel graphics hardware [10]. The main idea of these narrow-band approaches is to carry out the level-set process only in a small region around the level set of interest. Similar to active contours [9, 13], this requires the knowledge about the location of the zero level set. However, contrary to active contours, where the zero level set is explicitly deformed until it converges to the desired surface, the main idea of just deforming

the level-set function is maintained in the narrow-band methods. This allows for easy changes of the topology of the zero level set.

To our knowledge, there exists no method which combines the accuracy of directly applying level sets to unstructured data and the fast level-set computation using a narrow band. We propose such an approach that directly operates on the unstructured point-based data and does not need to reconstruct any scalar field at any positions other than the sample points.

3 GENERAL APPROACH

Let $f : D \rightarrow \mathbb{R}$ be a volumetric scalar field with bounded domain $D \subset \mathbb{R}^3$ given at a set of unstructured sample points $\mathbf{x}_i \in D$. Our goal is to efficiently extract a smooth isosurface $\Gamma_{\text{iso}} \subset D$ with respect to a given isovalue f_{iso} . This task should be carried out utilizing a narrow-band level-set approach which only operates on the sample points. The level-set approach provides a handle to control the smoothness of the resulting surface. The visualization pipeline consists of three main phases: the initialization, the level-set evolution, and the rendering. Its flow chart is depicted in Figure 1.

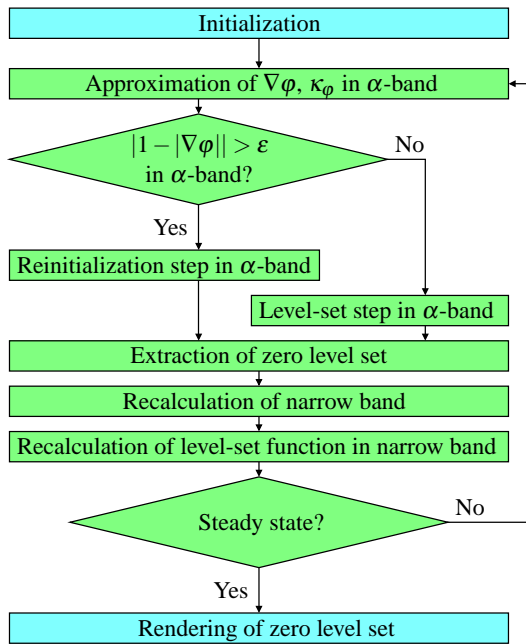


Figure 1: Illustration of the whole visualization pipeline. In the initialization phase, the level-set function and the narrow band are initialized. Next, the level-set function and the narrow band are consecutively updated following Equation (2) until the level-set function reaches steady state. Finally, the zero level set is rendered as resulting smooth isosurface.

In the initialization phase, we build a data structure that stores and handles neighborhood information, to efficiently compute level-set function properties, extract isopoints, and generate the narrow band. This is done using a three-dimensional kd -tree. In a preprocessing step, we compute the n nearest neighbors for

each sample. Here, $n = 26$ is chosen, inspired by the regular case on a structured equidistant hexahedral grid, where every sample \mathbf{x} has 26 nearest neighbors in the L_∞ -metric.

Afterwards, the level-set function $\varphi : D \rightarrow \mathbb{R}$ is initialized for every sample of the data set as a signed-distance function to a sphere. The points of the zero level set are extracted using direct isosurface extraction [18], which results in a point cloud representation of the extracted surface. Around the extracted isosurface, a narrow band of sample points consisting of two layers is generated. The inner layer or α -band consists of all sample points, which have a smaller distance to the isosurface than a given band width d_α . The outer layer or β -band consists of all sample points, which are not lying in the α -band but are needed to compute gradients or mean curvature at α -band points. The actual level-set process will be performed only within the α -band.

In the iterative processing phase, first the gradient $\nabla\varphi$ of the level-set function and the mean curvature κ_φ are approximated for the sample points in the α -band. The approximation is computed using a four-dimensional least-squares approach. If the norms of the level-set function gradients exceed a given threshold, a reinitialization step is applied to the level-set function with respect to the special Eikonal equation

$$\frac{\partial\varphi}{\partial t} = \text{sign}(\varphi)(1 - |\nabla\varphi|). \quad (1)$$

The reinitialization step brings the level-set function again close to a signed-distance function, which assures good numerical behavior of the whole process. If no reinitialization is necessary, a level-set step is performed in the α -band with respect to the equation

$$\frac{\partial\varphi}{\partial t} = ((1 - \lambda)(f - f_{\text{iso}} - \varphi) + \lambda\kappa_\varphi) |\nabla\varphi|, \quad (2)$$

which models a weighted combination of hyperbolic normal advection [15] and mean curvature flow [6] with smoothness parameter $\lambda \in [0, 1]$. Since we use an explicit Euler time discretization for updating the level-set function at the sample points, the time steps are bounded by the Courant-Friedrichs-Lewy (CFL) condition [3] to permit numerical stability. In particular, this condition also assures that the zero level set cannot leave the α -band in a single iteration.

After updating the function values in the α -band, the new zero level set is extracted and the narrow band has to be updated. Since we only update the function values of the points in the α -band, the points that have been added to the α -band after the level-set step do not hold the correct level-set function values. Consequently, the values have to be recomputed with respect to the new zero level set. This is done by interpolating between computed level-set function values and signed-distance

function values in the α -band. Similarly, the points in the β -band have not been updated and may hold wrong level-set function values. These points are updated by assigning signed-distance function values to the new zero level set.

The local level-set process is executed until the level-set function reaches steady state, i. e., until the function values do not change more than a given threshold from one time step to the subsequent one. After convergence, the zero level set has the desired properties of Γ_{iso} and is rendered using splat-based ray tracing [11].

4 NARROW BAND CREATION

Since the movement of the zero level set during the level-set iteration is bounded in each step by the CFL-condition, it is possible to restrict the level-set process to a band around the zero level set without losing flexibility or distorting the result. However, the level-set process at a sample point requires the approximation of level-set function derivatives in this point and, therefore, the level-set function values of its nearest neighbors. Restricting level-set update and derivative approximation to points within the band around the zero level set would, consequently, lead to oscillations at the border of the band, since derivatives there are less precisely approximated than near the zero level set.




Figure 2: Construction of the two-layer band around the zero level set (colored blue). The zero level set is extracted in form of a point cloud representation. Then, all sample points with distance to the zero level set less than d_α are marked as belonging to the inner layer of the band (green). Thereafter, all additional sample points needed for the gradient computations within the level-set process are marked as belonging to the outer layer of the band (red points).

We chose, instead, to use a band consisting of two nested layers. We want to use the notation α -band and β -band for the two layers. The construction of these layers is illustrated in Figure 2. The α -band describes the volume, where the actual level-set process is performed. All sample points having a smaller distance to the zero level set than the width d_α of the α -band are marked as members of the α -band. As the zero level set is extracted in a point cloud representation, the distance

of a sample point to the zero level set is defined as the distance to the nearest point of the zero level set. For each sample point the distance is efficiently obtained by a nearest-neighbor query to the kd -tree of the points of the zero level set.

The size of d_α depends on the distribution of the sample points, i. e., is data-dependent. To capture the entire area surrounding the zero level set, d_α has to be at least greater than the maximum distance of neighboring points within the point cloud representation of the zero level set. In fact, it should be even twice as large to permit a reasonable recalculation of the level-set function values. We discuss this in more detail in Section 5. In all our experiments, it was sufficient to choose one global value for d_α for the whole local level-set process, even for data sets with highly varying point densities. Hence, d_α was chosen depending on the sample distribution of the data set and independent of the zero level sets.

For the processing of the points in the α -band with respect to Equations (1) and (2), sample points outside the α -band are needed to approximate level-set function derivatives. Having established the α -band, all additional sample points that are required for derivative approximation of the points belonging to the α -band are marked as members of the β -band, which merely acts as support for the α -band processing.

5 NARROW BAND UPDATE

In the initialization phase, the level-set function is set globally to a radial signed-distance function. Hence, the level-set function values of points in the initial narrow band are directly given. However, after executing one local level-set step in the α -band, extracting the new zero level set, and updating the narrow band, the sample points may have changed their band membership. Consequently, some sample points may have been added to the α -band. For these new members of the α -band, the level-set function values are outdated.

Moreover, function values of all the points in the β -band have not been updated in the level-set step, but their new values are required for the subsequent level-set iteration. Hence, additional updates of the level-set function values at some sample points of the narrow band are likely to be required in each iteration step. An illustration of the idea is given in Figure 3.

For points with distance smaller than $\frac{d_\alpha}{2}$ to the zero level set, we can be sure that the level-set function values have been updated in the level-set iteration step. This is due to the fact that the level-set function value update is executed with a step size that is bounded by the CFL condition. Hence, no adjustment is necessary.

Sample points in the narrow band with distance to the zero level set larger than $\frac{d_\alpha}{2}$ (red points in Figure 3) may have not been updated or may even not have been in the narrow band at all in the preceding iteration step.




Figure 3: Updating the level-set function on the narrow band with size d_α : Points (green) with minimum distance to the zero level set points (blue) smaller than $\frac{d_\alpha}{4}$ have correctly been updated in the level-set iteration step. Points (red) in the narrow band with distance to the zero level set greater than $\frac{d_\alpha}{2}$ might have not been included in the computations of the last level-set iteration step. We assign their level-set function value to the signed distance to the zero level set points. For all points in the α -band lying in between, the new level-set function value is interpolated between the old level-set function value and their signed distance to the zero level set points.




Figure 4: Derivation of correct sign for signed-distance function generation at sample points with distance to the zero level set greater than $\frac{d_\alpha}{4}$. Each point of the zero level set \mathbf{x}_{iso} holds a surface normal \mathbf{n} , inducing an orientation of the zero level set and therewith inducing a well-defined sign for all points with distance greater than $\frac{d_\alpha}{4}$.

Hence, the level-set function value needs to be recomputed as illustrated in Figure 4. Since the level-set function is always kept near a signed-distance function it is feasible to reset the level-set function values in these points to a signed-distance function. Let \mathbf{x} be such a sample point that needs to be updated. Then, we make use of the - already computed - distance d to the nearest zero level set point \mathbf{x}_{iso} . The sign for the signed-distance function is directly induced by the surface normal of the nearest zero level set point as illustrated in Figure 4. Altogether, we set the level-set function value for \mathbf{x} to

$$\varphi_{\mathbf{x}} := \text{sign}(\langle \mathbf{n}, \mathbf{x} - \mathbf{x}_{\text{iso}} \rangle) |\mathbf{x} - \mathbf{x}_{\text{iso}}| .$$

Although the level-set function is kept near a signed-distance function at points with distance $d \leq \frac{d_\alpha}{2}$ to the zero level set and is explicitly been set to a signed-distance function at all the other points of the narrow band, there might occur a discontinuity at distance $d = \frac{d_\alpha}{2}$ to the zero level set. This can be avoided by

interpolating between level-set function value and constructed signed-distance function value in a transition zone. We chose this transition zone to be the distance range $d \in \left[\frac{d_\alpha}{4}, \frac{d_\alpha}{2} \right]$ to the zero level set. For a continuous transition, we apply an interpolation using the monotone polynomial

$$k : [0, 1] \rightarrow [0, 1], \quad k(t) := 6t^5 - 15t^4 + 10t^3 .$$

Consequently, for each sample point \mathbf{x} with distance $d \in \left[\frac{d_\alpha}{4}, \frac{d_\alpha}{2} \right]$ to the zero level set, we recalculate its level-set function value by

$$\varphi_{\mathbf{x}} := \varphi_{\mathbf{x}} \cdot \left(1 - k \left(\frac{4d - d_\alpha}{d_\alpha} \right) \right) + d \cdot k \left(\frac{4d - d_\alpha}{d_\alpha} \right),$$

which produces a C^2 -continuous blending between both functions.

In summary, keeping the level-set values for points with distance $d \in \left[0, \frac{d_\alpha}{4} \right)$, blending between level-set function and signed-distance function for points with distance $d \in \left[\frac{d_\alpha}{4}, \frac{d_\alpha}{2} \right]$, and setting all other points in the narrow band to a signed-distance function results in a continuous recalculation of the level-set function in the narrow band.

6 LOCAL LEVEL-SET PROCESS

To process the level-set function according to hyperbolic normal advection and mean curvature flow as modeled in Equation (2), we apply the methods already used in the global case [19] to estimate the level-set function gradient $\nabla \varphi$ and mean curvature κ_φ in each sample point in the α -band of each time step. Therefore, only nearest-neighbor computations for the sample points are used. Furthermore, no information about the scalar field other than at the sample points in the narrow band is needed.

Since the quality of the level-set process significantly degrades if the level-set function φ is not close to a signed-distance function of the zero level set, the level-set function is initialized as a signed-distance function and the function values of points in the β -band get reset to represent a signed-distance function to the zero level set. However, the level-set process cannot maintain this property. To keep the level-set function near a signed-distance function, it is reinitialized using a PDE-based approach solving the special Eikonal equation (1), when the norms of the gradients of the function exceed a certain threshold.

Since the level-set processes following Equations (1) and (2) are performed using an explicit time discretization of order one, the used time step has to be chosen carefully to keep the zero level set always in the α -band and fulfill the CFL-condition. In fact, meeting the CFL-condition is also sufficient to keep the zero level set within the α -band during computations, as already observed by Peng et al. [17].

7 RESULTS AND DISCUSSION

The presented approach was applied to a variety of data sets to verify our method and evaluate it in terms of accuracy and speed. For performance analysis we applied it to an unstructured point-based volume data set with eight million randomly distributed samples. The data set was generated by resampling the regular Hydrogen data set of size $128 \times 128 \times 128$ to the random points. An illustration of the evolution process for this data set is shown in Figure 5.




Figure 5: Evolution of the zero level set when applying the local level-set method with $\lambda = 0.01$ to the Hydrogen data set with eight million sample points. For each time step, a splat-based ray tracing of the zero level set is shown on the right-hand side. On the left-hand side, a point rendering of a slab of the data set is shown illustrating the narrow band. Extracted surface points of the zero level set are colored black, sample points in the α -band are colored green, and sample points in the β -band are colored red. Sample points not belonging to the narrow band are not rendered.

The same data set was used to analyze the performance of our method in terms of speed. The data set was downsampled to different sizes for investigating the scalability of our approach. All computation times were measured on a single 2.66 GHz XEON processor.

The runtime analysis for the preprocessing phase is given in Table 1. We observe that most of the time is required for the nearest neighbor computation. However, the data to which we typically apply our method stem from numerical simulations, where nearest neighbors are known such that this pre-processing step can be dropped.

samples	<i>kd</i> -tree gen.	NN calc.	<i>kd</i> -neig. calc.
2M	4 sec	111 sec	4.1 sec
4M	9 sec	239 sec	10.1 sec
8M	19 sec	541 sec	21.9 sec

Table 1: Computation times for the preprocessing of the Hydrogen data set with different sample quantities, including the generation of the *kd*-tree, calculation of 26 nearest neighbors (only required if not given), and the calculation of the *kd*-tree neighbors needed for zero level-set surface extraction.

The computation times for one level-set iteration step analyzed with respect to the Hydrogen data set are shown in Table 2. It comprises a summary of the results including the number of samples, the number of extracted zero level set points, and the calculation times of all computation steps, i. e., for the zero level set extraction, for the update of the narrow band (including narrow band function value updates), and for the level-set process.

samples	points	point ex.	band up.	ls step
2M	18k	1.9 sec	1.8 sec	1.5 sec
4M	29k	2.6 sec	2.4 sec	2.8 sec
8M	47k	4.1 sec	3.9 sec	5.2 sec

Table 2: Computation times for one level-set iteration step of the Hydrogen data set with different numbers of sample points. The number of extracted zero level set points as well as the computation times for zero level set extraction are given. Moreover, the computation times for the narrow band update and the calculation of one level-set step are given.

The whole local level-set process for extracting a smooth isosurface from the Hydrogen data set with eight million sample points and given nearest neighbors needed 24 steps and was performed in 6 minutes. For the four million version of the data set, the overall computation time for the entire level-set approach including pre-computations dropped to 84 seconds. This is a significant speed-up in comparison to the time of 13 minutes for the global level-set process [19] for the four million Hydrogen data set.¹ When comparing the computation times per iteration (including computation of zero level set, narrow band, and level-set function

¹ These numbers assume that the nearest neighbors are known.

update), the 7.8 seconds of our presented local level-set method are one order of magnitude faster than the 67.8 seconds for the global level-set method [19]. Still, our local level-set method produces equivalent results in terms of quality and correctness.

Next, we have analyzed the speed of the iterative phase of our approach with respect to the number of points in the α -band. We applied the method to a re-sampled unstructured point-based data set of eight million randomly distributed samples, generated from the regular Engine data set of size $256 \times 256 \times 128$. A rendering of the zero level set after convergence is shown in Figure 6.




Figure 6: Splat-based ray tracing of the zero level set extracted after convergence of the local level-set method with $\lambda = 0.01$. The method was applied to the Engine data set with eight million unstructured sample points. (Data set courtesy of General Electric.)

Just by looking at the rendering, the complex topology and geometry of the zero level set is obvious. Consequently, the number of isopoints steadily increases during the local level-set evolution as well as the number of sample points in the narrow band. A comparison of the computation speed of our local method with different number of extracted zero level set points is given in Table 3. Note that the number of points in the α -band is increasing sub-linearly with the number of extracted zero level set points, since several surface parts get very close to each other and their neighboring regions share data points in the band.

ls points	α -points	point ex.	band u.	ls step
20k	130k	2.1 sec	2.0 sec	3.5 sec
49k	287k	4.3 sec	3.8 sec	6.3 sec
83k	404k	7.6 sec	6.9 sec	8.5 sec

Table 3: Computation times for the iteration phase of the Engine data set at different steps of the level-set process with different numbers of zero level set points. For each step, the number of extracted zero level set points, the number of points in the α -band, and all computation times within an iteration, i. e., for extracting the zero level set, updating the narrow band, and performing one level-set step, are given.

Finally, we applied our local level-set method to unstructured point-based volume data from a real scien-

tific application. The data set was provided by astrophysical particle simulations [20] of Stephan Rosswog, Jacobs University, Bremen, Germany. In the simulation, a set of particles representing a White Dwarf is passing a black hole and is torn apart by the strong gravity. The data set represents one time step of this simulation and consists of 500k sample points. The density of the points in space is varying in two orders of magnitude, directly showing that scattered data interpolation techniques would not be applicable due to the introduction of enormous errors.

We applied our local level-set method to this practical data set to compare the results of the local and the global method. Both methods were applied with the same initial signed-distance function, the same iso-value, and the same smoothness parameter. Renderings of the obtained zero level sets after convergence of both methods are shown in Figure 7. The presented local level-set method required only 6 minutes of computation time (including all computations) until reaching steady state. In comparison, the global level-set process using actually asynchronous time integration required 68 minutes. In terms of quality, there are no significant differences visible between the resulting surfaces. We have measured the distance between the points of the two extracted zero level sets. The maximum deviation between the surfaces was always of the same order as the stopping criterion for the level-set process. This results from the linear interpolation of surface points from the quasi-linear level-set function.

All experiments have shown, that the proposed local level-set method is significantly faster than the global method but equivalent in terms of quality. The narrow-band method is also able to process objects with complex topology and is applicable to real-world data.

8 CONCLUSION

We have presented a local level-set method that combines the accuracy of directly applying level sets to unstructured data and the fast level-set computation using narrow bands. We directly operate on the unstructured data and do not need to process data at any positions other than the sample points, effectively avoiding interpolation errors. For each iteration, we extract the zero level set, build a narrow band of sample points around it and process the level-set function only in the narrow band with respect to a combination of normal advection with mean curvature flow. The final zero level set is visualized using point-based rendering.

Our method is capable of robustly processing data sets with several million sample points and highly varying point density. For all investigated examples we achieved a speed-up of one order of magnitude compared to the global level-set method. Nevertheless, the proposed method achieves the same results in terms of quality.



Figure 7: Comparison of global and local level sets for the White Dwarf data set with 500k unstructured sample points with highly varying point density. Both methods have been applied to the same initial signed-distance function using the same isovalue and the same smoothness parameter $\lambda = 0.1$. The zero level set obtained by the global level-set method with asynchronous time integration is shown on the left-hand side. In comparison, the smooth isosurface generated by our local level-set method, shown on the right-hand side, exhibits no significant difference in terms of quality, but was generated more than ten times faster.

ACKNOWLEDGMENTS

This work was supported by the Deutsche Forschungsgemeinschaft (DFG) under project grant LI-1530/6-1.

REFERENCES

- [1] David Adalsteinsson and James A. Sethian. A fast level set method for propagating interfaces. *Journal of Computational Physics*, 118(2):269–277, 1995.
- [2] David Breen, Ross Whitaker, Ken Museth, and Leonid Zhukov. Level set segmentation of biological volume datasets. In Jasjit S. Suri, David L. Wilson, and Swamy Laxminarayan, editors, *Handbook of Biomedical Image Analysis, Volume I: Segmentation Models, Part A*, pages 415–478, New York, NY, USA, 2005. Kluwer.
- [3] Richard Courant, Kurt Otto Friedrichs, and Hans Lewy. On the partial difference equations of mathematical physics. *IBM Journal of Research and Development*, 11(2):215–234, 1967.
- [4] Herbet Edelsbrunner and Nimish R. Shah. Incremental topological flipping works for regular triangulations. In *SCG '92: Proceedings of the eighth annual symposium on Computational geometry*, pages 43–52, New York, NY, USA, 1992. ACM.
- [5] Douglas Enright, Frank Losasso, and Ronald Fedkiw. A fast and accurate semi-lagrangian particle level set method. *Computers and Structures*, 83(6–7):479–490, 2004.
- [6] Lawrence C. Evans and Joel Spruck. Motion of level sets by mean curvature I. *Journal of Differential Geometry*, 33(3):635–681, 1991.
- [7] Richard Franke and Gregory M. Nielson. Scattered data interpolation: A tutorial and survey. In Hans Hagen and Dieter Roller, editors, *Geometric Modeling: Methods and Applications*, pages 131–160. Springer, New York, NY, USA, 1991.
- [8] Simone E. Hieber and Petros Koumoutsakos. A lagrangian particle level set method. *Journal of Computational Physics*, 210(1):342–367, 2005.
- [9] Hon Pong Ho, Yunmei Chen, Huafeng Liu, and Pengcheng Shi. Level set active contours on unstructured point cloud. In *CVPR '05: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 2, pages 690–697, Los Alamitos, CA, USA, 2005. IEEE Computer Society.
- [10] Aaron E. Lefohn, Joe M. Kniss, Charles D. Hansen, and Ross T. Whitaker. A streaming narrow-band algorithm: Interactive computation and visualization of level sets. *IEEE Transactions on Visualization and Computer Graphics*, 10(4):422–433, 2004.
- [11] Lars Linsen, Karsten Müller, and Paul Rosenthal. Splat-based ray tracing of point clouds. *Journal of WSCG*, 15:51–58, 2007.
- [12] Chohong Min. Local level set method in high dimension and codimension. *Journal of Computational Physics*, 200(1):368–382, 2004.
- [13] Bryan S. Morse, Weiming Liu, Terry S. Yoo, and Kalpathi Subramanian. Active contours using a constraint-based implicit representation. In *CVPR '05: Proceedings of the 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05) - Volume 1*, pages 285–292, Washington, DC, USA, 2005. IEEE Computer Society.
- [14] Ken Museth, David E. Breen, Leonid Zhukov, and Ross T. Whitaker. Level set segmentation from multiple non-uniform volume datasets. In *VIS '02: Proceedings of the conference on Visualization '02*, pages 179–186, Washington, DC, USA, 2002. IEEE Computer Society.
- [15] Stanley Osher and Ronald Fedkiw. *Level set methods and dynamic implicit surfaces*. Springer, New York, NY, USA, 2003.
- [16] Stanley. Osher and James A. Sethian. Fronts propagating with curvature-dependent speed: Algorithms based on Hamilton-Jacobi formulations. *Journal of Computational Physics*, 79(1):12–49, 1988.
- [17] Danping Peng, Barry Merriman, Stanley Osher, Hongkai Zhao, and Myungjoo Kang. A PDE-based fast local level set method. *Journal of Computational Physics*, 155(2):410–438, 1999.
- [18] Paul Rosenthal and Lars Linsen. Direct isosurface extraction from scattered volume data. In Beatriz Sousa Santos, Thomas Ertl, and Kenneth I. Joy, editors, *EuroVis06: Proceedings of the Eurographics/IEEE-VGTC Symposium on Visualization*, pages 99–106, Aire-la-Ville, Switzerland, 2006. Eurographics Association.
- [19] Paul Rosenthal and Lars Linsen. Smooth surface extraction from unstructured point-based volume data using PDEs. *IEEE Transactions on Visualization and Computer Graphics*, 14(6):1531–1546, 2008.
- [20] Stephan Rosswog, Enrico Ramirez-Ruiz, William Raphael Hix, and Marius Dan. Simulating black hole white dwarf encounters. *Computer Physics Communications*, 179(1–3):184–189, 2008.
- [21] James A. Sethian. *Level Set Methods and Fast Marching Methods*. Cambridge University Press, Cambridge, UK, 2nd edition, 1999.
- [22] Mark Sussman and Emad Fatemi. An efficient, interface-preserving level set redistancing algorithm and its application to interfacial incompressible fluid flow. *SIAM Journal on Scientific Computing*, 20(4):1165–1191, 1999.