

Efficient GPU Fluid Visualization

Dávid Illés, László Szécsi

Abstract

In this paper we describe a method for rendering particle-based fluid representations, without screen-space filtering approximation. The algorithm builds per-pixel lists of relevant particles, then incrementally constructs a piecewise polynomial approximation of summed base functions along rays, and finds intersections with the isosurface using those. This new approach scales well for a high number of particles, it can handle local extremities of depth complexity robustly, and it does not suffer from the inaccuracies and limitations of screen-space approximation methods.

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [I.3.7]: Three-Dimensional Graphics and Realism: Raytracing

1. Introduction

Lagrangian base functions^{5,1} are widely used to reconstruct fluid density from particle-based simulations, from hydrodynamics to molecular modeling. Such simulations, even with hundreds of thousands of particles, are becoming available in real-time applications, where a fast but general visualization method is needed.

Each particle has a base density function, and a set of particles represents a smooth surface as the isosurface of the density field (see Figure 2 for renderings of isosurfaces generated by two particles, and Figure 3 for the effect of different effective radii). The isosurface is mainly visualized by polygonization (the marching cubes algorithm⁴), ray casting⁶, or approximated by screen-space filtering of a solid estimating surface¹⁰. The marching cubes algorithm can suffer from voxelization artifacts for low resolution grids, or be extremely computation and memory intensive when using high resolution. Ray casting produces high quality smooth surfaces, but the ray-isosurface intersection test is also very expensive to perform. Finally, screen-space filtering can produce visually acceptable results in real-time under controlled conditions, but it is not useful when close-ups or other exact details are required.

We propose a fast ray casting method to render soft fluid particles on the GPU (Figure 1). The approach is similar to that of Kanamori et al.³, gathering particles along rays, and maintaining relevant density information as we process them. However, we replace depth peeling⁸ (which is multi-pass, and allows limited depth complexity) with gathering

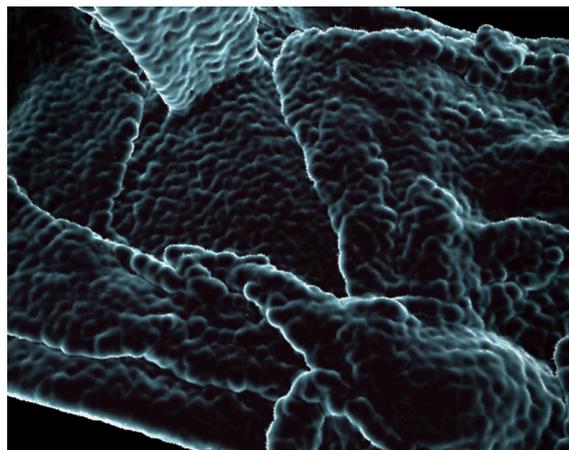


Figure 1: Close-up of a scene with 100,000 particles rendered using cel shading at 13 FPS.

per-pixel fragment lists in a single pass. Also, we use a piecewise cubic approximation of the density function, which allows for faster intersection computation, replacing Bezier clipping. These points also make the main contributions of the paper. Otherwise, all considerations by Kanamori et al. concerning related work apply. The resulting method scales well with the number of particles, allowing real-time visualization of massive particle simulations.



Figure 2: Two particles rendered with different isosurface levels and environment mapping.

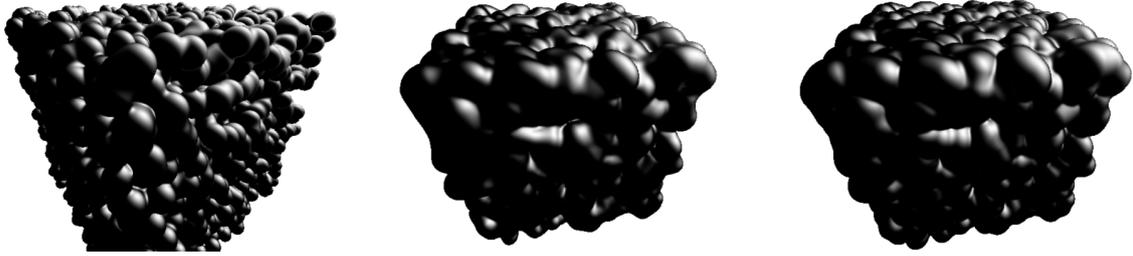


Figure 3: Pitch renderings with different effective radii.

2. The Lagrangian model

The *density function* ρ_j for particle j is a function of r , the distance from the particle center \mathbf{p}_j , and monotonically decreases with r . For M particles, the shape of the curved surface at density value L is defined by the points \mathbf{x} satisfying

$$\rho(\mathbf{x}) = \sum_{j=0}^{M-1} \rho_j(\|\mathbf{x} - \mathbf{p}_j\|) - L = 0, \quad (1)$$

where $\rho(\mathbf{x})$ is called the *total density*. Although Gaussian and hyperbolic kernel density functions were also used, they are impractical for computations as all terms of the summation in Equation 1 are non-zero. Instead, ρ_j is usually defined to have finite support R_j , with $\rho_j(r) = 0$ where $r \geq R_j$. Within $r < R_j$, $\rho_j(r)$ is typically given in polynomial form, which is not only easy to evaluate, but it can be chosen to meet continuity criteria. One of the most widely used density functions is the elegant fifth-degree version by Perlin⁷

$$\rho_j^{\text{Perlin}}(r) = \begin{cases} 0 & \text{if } r \geq R_j, \\ 1 - \frac{r^3}{R_j^3} \left(\frac{r}{R_j} \left(6 \frac{r}{R_j} - 15 \right) + 10 \right) & \text{if } r < R_j, \end{cases}$$

which ensures tangential continuity. The six-degree polynomial proposed by Wyvill et al.¹¹

$$\rho_j^{\text{Wyvill}}(r) = \begin{cases} 0 & \text{if } r \geq R_j, \\ 1 - \frac{4}{9} \frac{r^6}{R_j^6} + \frac{17}{9} \frac{r^4}{R_j^4} - \frac{22}{9} \frac{r^2}{R_j^2} & \text{if } r < R_j, \end{cases}$$

can be written as a function of r^2 , meaning that taking a square root for the vector norm computation in Equation 1 is not necessary, and $\rho(\mathbf{x})$ itself is piecewise polynomial.

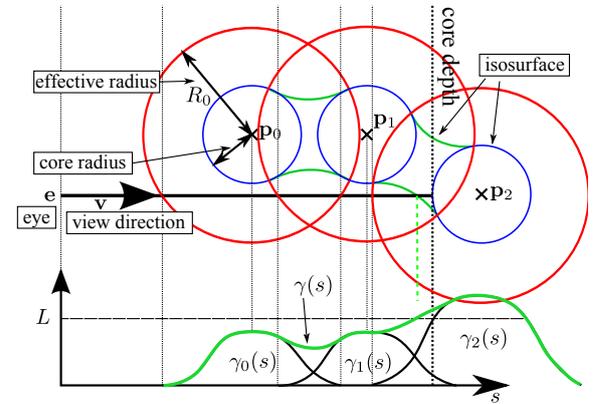


Figure 4: Nomenclature for particles and their density functions along a ray.

Figure 4 explains the nomenclature. We call R_j the *effective radius* of particle j , and the sphere of centre \mathbf{p}_j and radius R_j is the *effective sphere*. The radius for which the density function equals L is the *core radius*. The *core sphere* of this radius is always contained within the isosurface.

If we wish to visualize this implicit surface using ray casting, we need to be able to find the intersection of the isosurface and a ray. Let the ray equation be

$$\mathbf{y}(s) = \mathbf{e} + s\mathbf{v},$$

where \mathbf{e} is the eye position and \mathbf{v} is the viewing direction. Substituting this into Equation 1 gives:

$$\rho(\mathbf{y}(s)) = \sum_{j=0}^{M-1} \rho_j(\|\mathbf{e} + s\mathbf{v} - \mathbf{p}_j\|) - L.$$

Let us introduce $\gamma(s) = \rho(\mathbf{y}(s))$ for the total density as a function of the ray parameter. Let $\gamma_j(s)$ be the j th term in the sum, which is an individual kernel density, also as the function of the ray parameter.

3. Polynomial approximation

If ρ_j can be written as a function of r^2 , $\gamma_j(s)$ will be piecewise polynomial. The fact that such a choice of a density function is possible is an important motivation for using a piecewise polynomial approximation $\tau(s)$ for $\gamma(s)$. We construct $\tau(s)$ as a sum of $\tau_j(s)$ individual piecewise polynomial approximations of $\gamma_j(s)$ particle density functions. Further on in this discussion, we are going to assume piecewise cubic $\tau_j(s)$, and we provide a method for their construction from density kernel functions $\rho_j(r)$. However, it is easy to extend our method to higher degrees, and it is also possible to use exact polynomials for appropriately chosen density functions. The use of cubics introduces some innoticeable view-dependent distortion, but it allows for simpler implementation and faster evaluation.

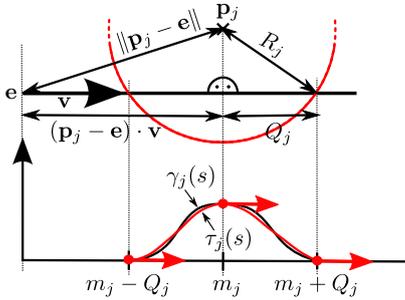


Figure 5: Effective interval computation and third order spline approximation of the density function.

The original $\gamma_j(s)$ is non-zero where the ray intersects its effective sphere (see Figure 5). For the *midpoint* of the intersected interval, the ray parameter m_j can be computed as

$$m_j = (\mathbf{p}_j - \mathbf{e}) \cdot \mathbf{v},$$

and the half-length of the intersected segment (using the Pythagorean theorem) is

$$Q_j = \sqrt{R_j^2 - \left(\|\mathbf{p}_j - \mathbf{e}\|^2 - m_j^2\right)}.$$

We call the interval $[m_j - Q_j, m_j + Q_j]$ the *effective interval*. For all practical, smooth density functions, the values and derivatives of $\gamma_j(s)$ are zero at the endpoints of the effective interval (also shown in Figure 5). The derivative is also zero at the midpoint. We construct $\tau_j(s)$ by fitting a third-order spline on these characteristic points and tangents. This results in two cubic segments in $[m_j - Q_j, m_j]$ and $[m_j, m_j + Q_j]$. Thus, with \mathbf{s} as $(s^3, s^2, s, 1)$,

$$\tau_j(s) = \begin{cases} 0 & \text{if } s \notin [m_j - Q_j, m_j + Q_j], \\ \mathbf{s} \cdot \mathbf{u}_j^\uparrow & \text{if } s \in [m_j - Q_j, m_j], \\ \mathbf{s} \cdot \mathbf{u}_j^\downarrow & \text{if } s \in [m_j, m_j + Q_j], \end{cases}$$

where the coefficients can be computed by solving the system of equations written for the values and the derivatives:

$$\tau_j(m_j - Q_j) = \tau_j(m_j + Q_j) = 0,$$

$$\tau_j(m_j) = \gamma_j(m_j),$$

$$\tau_j'(m_j - Q_j) = \tau_j'(m_j) = \tau_j'(m_j + Q_j) = 0,$$

yielding

$$\mathbf{u}_j^\uparrow = \frac{\gamma_j(m_j)}{Q_j^3} \begin{pmatrix} -2 \\ 3(2m_j + Q_j) \\ -6(m_j^2 + m_j Q_j) \\ 2m_j^3 + 3Q_j m_j^2 \end{pmatrix}, \quad (2)$$

$$\mathbf{u}_j^\downarrow = \frac{-\gamma_j(m_j)}{Q_j^3} \begin{pmatrix} -2 \\ 3(2m_j + Q_j) \\ -6(m_j^2 + m_j Q_j) \\ 2m_j^3 + 3Q_j m_j^2 - Q_j^3 \end{pmatrix}. \quad (3)$$

4. Ray decomposition

As $\tau(s)$ is the sum of piecewise cubics, it is piecewise cubic itself, with all end- and midpoints of effective intervals as *subdomain separators*. Let us denote these points in increasing order by s_i , with $i = 1 \dots n$, where n is the number of such points, and prepend $s_0 = 0$ to this ordered list.

$$\gamma(s) \approx \tau(s) = \mathbf{s} \cdot \mathbf{w}_i \quad \text{if } s_i \leq s < s_{i+1} \quad \Big| \quad i \in \{0, \dots, n-1\}$$

where \mathbf{w}_i are the polynomial coefficient vectors. In every subdomain, the density function can be derived from a given set of particles. Our goal is to identify all these subdomains, compute the approximating polynomials, and solve it for intersection with the isosurface.

Along the ray, the coefficients of $\tau(s)$ change at every subdomain separator. The change of coefficients $\Delta \mathbf{w}_i$ at s_i can be

deduced from the cubic coefficients as

$$\Delta \mathbf{w}_i = \begin{cases} \frac{\gamma_j(s_i)}{Q_j^3} \begin{pmatrix} -2 \\ 6s_i + 3Q_j \\ -6s_i^2 - 6s_i Q_j \\ 2s_i^3 + 3s_i^2 Q_j \end{pmatrix} & \text{if } s_i = m_j - Q_j, \\ \frac{\gamma_j(s_i)}{Q_j^3} \begin{pmatrix} 4 \\ -12s_i \\ 12s_i^2 \\ -4s_i^3 \end{pmatrix} & \text{if } s_i = m_j, \\ \frac{\gamma_j(s_i)}{Q_j^3} \begin{pmatrix} -2 \\ 6s_i - 3Q_j \\ -6s_i^2 + 6s_i Q_j \\ 2s_i^3 - 3s_i^2 Q_j \end{pmatrix} & \text{if } s_i = m_j + Q_j. \end{cases}$$

If the eye is not within the effective sphere of any particle, $\mathbf{w}_0 = \mathbf{0}$. Otherwise, the cubics of the intervals that contain the eye have to be summed. Coefficients for consequent subdomains can simply be found by applying the changes

$$\mathbf{w}_i = \mathbf{w}_{i-1} + \Delta \mathbf{w}_i.$$

5. The proposed method

Gathering a list of fragments for image pixels is possible with Shader Model 5 hardware, as it has been demonstrated in the Order-Independent Transparency (OIT) method². In the OIT method the fragments are gathered so that they can be sorted by depth to evaluate alpha-blending-like transparency. We use similar fragment lists to process intersected particles.

The algorithm consists of the following phases:

Sorting: We sort particles according to their *effective distance* ($\|\mathbf{p}_j - \mathbf{e}\| - R_j$) from the camera, in ascending order.

Lay down depth: We render solid scene and particle core depth into a *core depth* texture. The geometry rendered for core spheres is a covering billboard⁹ or some more fitting polygonal enclosing object (we did not find a significant performance difference). The pixel shader discards pixels in which the sphere is not intersected. Z-buffering is used, cores are rendered front-to-back for performance.

Gathering: We render the effective spheres of particles, gathering their IDs into per pixel linked lists, maintained in a read/write GPU buffer. Effective spheres are rendered with the same technique as core spheres in the previous phase. The intersected effective interval is clipped against the core depth. Finally, the particle ID is prepended to the linked list that belongs to the pixel. Particles are rendered back-to-front, thus the created lists are ordered by increasing effective distance.

Ray casting: We render a full-viewport quad. For every pixel, the lists are expanded to end- and midpoint records, which are sorted and used to evaluate ray–isosurface intersection. This process is detailed in Section 6.

Shading: The normal of the isosurface at the intersection point is computed, and the shading formula is evaluated to get the pixel color. The normal vector at \mathbf{x} can be derived from $-\nabla \rho(\mathbf{x})$.

6. Ray casting

In every pixel, we need to assemble an ordered list of end- and midpoints of effective intervals of all intersected particles, storing distances s_i and coefficient changes $\Delta \mathbf{w}_i$. The fact that particles are already sorted by effective depth helps minimize local shader memory usage and sorting overhead, because it allows us to process partial subdomain lists. For every pixel, the shader performs the following steps:

1. The polynomial coefficient vector \mathbf{w} , which is a running variable, is initialized to zero.
2. For every particle in the list:
 - a. The effective sphere is intersected with the ray through the pixel. Intersected intervals are clipped to between zero and the core depth. Separator records $(s_i, \Delta \mathbf{w}_i)$ for end- and midpoints are inserted into a list ordered by depth. This list is short and it is stored in a fixed-sized local-memory array. The sorting algorithm needs to work well with relatively few — surely less than 100 — fragments. Thus, computational complexity is trumped by less storage and administration overhead, making insertion sort the optimal choice.
 - b. *Safe* separator records are those separators in the list, which are within the effective distance of the next, yet unprocessed particle (in the *safe zone*). As particles are ordered, separators of further effective spheres cannot precede the safe ones. For these safe separators, the polynomial coefficients are updated, and the intersection within subdomains is evaluated. The polynomial coefficients are maintained in the running variable \mathbf{w} , adding $\Delta \mathbf{w}_i$ as the separators are processed. The intersection computation only requires the solution of a cubic, where the practically exclusive monotonous case can easily be identified and very effectively handled. The processed separators are discarded from the ordered list. If no particles are left, we can evaluate all remaining subdomains.
3. When an intersection was found, we compute the normal and shading.

We store the elements of the ordered list in a local-memory array in reverse order. Elements are removed simply by decreasing the element count. For every intersected sphere we need to insert three records in known order, meaning records of a lesser distance have to be moved only once.

The fixed-size local memory array may be insufficient to hold separators from too many overlapping spheres, forcing us to process non-safe separators. This creates a trade-off between performance and the amount of overlapping we can



Figure 6: Particle formations rendered with environment mapping.

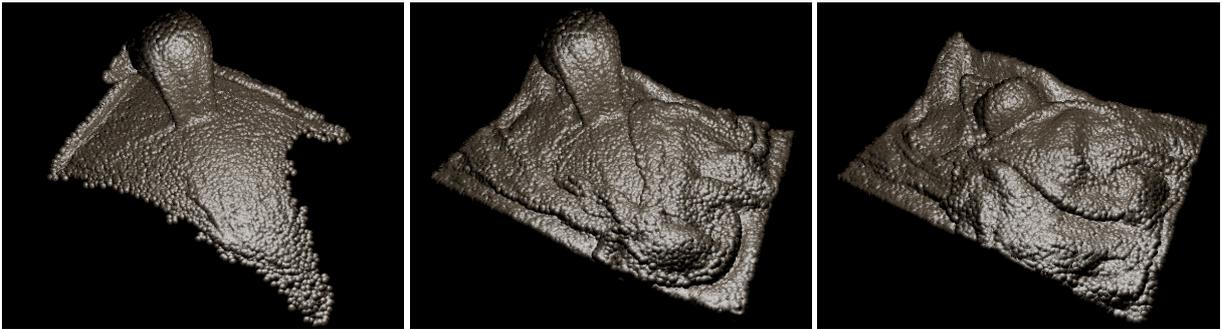


Figure 7: Frames of a mud simulation rendered as 200,000 particles. Frame rates are 25 FPS, 13 FPS, and 6 FPS.

handle robustly. Note that there is no limit on screen-space depth complexity, merely on the number of effective spheres a world-space point can belong to. As intersection records can be kept small, bounds in typical particle simulation scenarios are easily met.

7. Results and conclusion

Our approach was implemented in C++, using DirectX 11 and HLSL shaders (see Figure 6 for some screenshots). We ran our tests on a PC with an Intel Quad 2.40 GHz CPU and an NVIDIA GeForce GTX 560 Ti graphics card. According to our measurements the method runs real-time even for tens of thousands of particles (Table 7, Figure 7), and it can achieve interactive framerates for 1,600,000 particles. Our test scenes were simulated in *Next Limit's RealFlow*.

The cost of the ray casting phase, which used to be the bottleneck in earlier algorithms, becomes fairly constant after a number of particles. The $\mathcal{O}(M)$ complexity of rendering all particles and the $\mathcal{O}(M \log M)$ cost of sorting take over at large numbers. We can conclude that our approach scales well for massive amounts of particles, with parallelized sorting and *Potentially Visible Set* solutions being the broadest avenues for further improvement. We also wish to investigate applicability to computing multiple intersections along rays, allowing for translucent shading.

Particles	Total	Sort	Depth	Gather	Cast
50,000	40	8	2	6	24
100,000	77	15	5	15	42
200,000	160	32	12	36	80
800,000	505	165	65	191	84
1,600,000	991	350	131	402	108

Table 1: Total frame rendering and phase execution times in msec at 1024×768 . Sort includes CPU-GPU data transfer, Cast includes shading.

8. Acknowledgements

This work was supported by OTKA K-719922 and 101527, as well as the project of the “Development of quality-oriented and harmonized R+D+I strategy and functional model at BME” (Project ID: TÁMOP-4.2.1/B-09/1/KMR-2010-0002).

References

1. J.F. Blinn. A generalization of algebraic surface drawing. *ACM Transactions on Graphics (TOG)*, 1(3):235–256, 1982.
2. C. Everitt. Interactive order-independent transparency. *White paper, nVIDIA*, 2(6):7, 2001.
3. Y. Kanamori, Z. Szego, and T. Nishita. GPU-based fast ray casting for a large number of metaballs. In *Computer Graphics Forum*, volume 27, pages 351–360, 2008.
4. W.E. Lorensen and H.E. Cline. Marching cubes: A high resolution 3d surface construction algorithm. In *ACM Siggraph Computer Graphics*, volume 21, pages 163–169. ACM, 1987.
5. H. Nishimura, M. Hirai, T. Kawai, T. Kawata, I. Shirakawa, and K. Omura. Object modeling by distribution function and a method of image generation. *The Transactions of the Institute of Electronics and Communication Engineers of Japan*, 68(Part 4):718–725, 1985.
6. T. Nishita and E. Nakamae. A method for displaying metaballs by using bézier clipping. In *Computer Graphics Forum*, volume 13, pages 271–280. Wiley Online Library, 1994.
7. K. Perlin. Improving noise. In *ACM Transactions on Graphics (TOG)*, volume 21, pages 681–682. ACM, 2002.
8. L. Szirmay-Kalos and W. Purgathofer. Global ray-bundle tracing with hardware acceleration. In *Rendering Techniques '98*, pages 247–258, 1998.
9. T. Umenhoffer, L. Szirmay-Kalos, and G. Szijártó. Spherical billboards and their application to rendering explosions. In *Graphics Interface*, pages 57–64, 2006.
10. W.J. van der Laan, S. Green, and M. Sainz. Screen space fluid rendering with curvature flow. In *Proceedings of the 2009 symposium on Interactive 3D graphics and games*, pages 91–98. ACM, 2009.
11. G. Wyvill, C. McPheeters, and B. Wyvill. Data structure for soft objects. *The visual computer*, 2(4):227–234, 1986.