

Screen Space Foam Rendering

Nadir Akinci Alexander Dippel Gizem Akinci Matthias Teschner
nakinci,dippela,gakinci,teschner@informatik.uni-freiburg.de
University of Freiburg
Georges Koehler Allee 052
79110 Freiburg Germany

ABSTRACT

We present a method for the efficient rendering of large scale particle-based foam data in screen space using a GPU based rendering pipeline. Our approach employs a multi-pass rendering technique to imitate some of the effects that are commonly accomplished by using expensive ray-tracing based methods. We demonstrate through different scenarios that our pipeline is able to produce convincing foam renderings for large scale scenarios and it has a significant performance advantage compared to using ray-casting techniques for rendering such particle data.

Keywords

Rendering, Fluids, Foam, Particles

1 INTRODUCTION

1 Foam is a complex phenomenon whose behavior and
2 appearance is challenging to simulate in computer
3 graphics. When viewed from a close distance, foam is
4 composed of many air bubbles sticking to each other.
5 It can occur inside most fluids as a result of trapped
6 air. One can observe milky white foam caused by
7 dashing waves on seashores. For most semi-transparent
8 materials, it is an interesting observation that, even
9 though the underlying material may have a color, the
10 foam usually looks whitish to the viewer. The reason
11 for this behavior is that the foam is composed of thin
12 films of fluid containing air. As the number of such
13 thin films increase per unit volume, all incoming light
14 is reflected without allowing any light to penetrate
15 beneath it. This optical phenomenon makes the foam
16 look brighter than the material itself, to the point that it
17 looks almost white. This paper focuses on the efficient
18 rendering of such white foam by approximating some
19 important effects in screen space, that are otherwise
20 time consuming to compute in a physically correct
21 way. Our technique is specifically useful for complex
22 large-scale scenarios, where large amount of foam data
23 need to be rendered. In the remainder of this section,
24 we first summarize the existing works about GPU
25 accelerated rendering of fluid data (Sec. 1.1), foam

simulation and rendering (Sec.1.2) and then highlight
our contribution (Sec. 1.3).

1.1 GPU Rendering of Fluids

For non-interactive applications, fluid surfaces are
generally visualized by triangulating the isosurface of
the particle data (e.g. [ZB05, YT10, AIAT12]) and
then rendering the resultant mesh using ray-tracing
based techniques to produce convincing results. For
real-time applications, the computational overhead
of those approaches remains too high. Therefore,
for the efficient GPU accelerated visualization of
fluid surfaces, several methods have been proposed
in the recent years, e.g., using screen space surface
construction [MSD07, FAW10], height field tech-
niques [CM10] and methods that are based on particle
splatting [vdLGS09, BSW10]. Even though foam is
actually composed of the molecules of the underlying
fluid, its characteristic appearance requires it to be
handled using different rendering approaches, which
will be explained in the next section.

1.2 Foam Simulation and Rendering

In computer graphics, foam generation techniques are
used to enhance the realism of existing fluid simula-
tions. High quality foam simulation and rendering tech-
niques are commonly encountered in movies [GLR⁺06,
BSK⁺07] and in commercial fluid simulation and visu-
alization packages [hyb11]. In those works, however,
the underlying foam generation and rendering stages
are usually described briefly. Although foam is com-
posed of fluid and air mixture, some of the existing re-
search also focus on generating foam particles, usually
in a scale smaller than the fluid particles to be able to

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, or re-
publish, to post on servers or to redistribute to lists, requires
prior specific permission and/or a fee.



Figure 1: A flood scenario. Foam is rendered using our technique and composited with the rest of the scene (left and middle). Picture of real sea foam caused by a whirlpool (right) (©Reuters).

58 enhance the flow detail [TFK⁺03, GLR⁺06, LTKF08,
59 MMS09, IAAT12].

60 For high quality foam renderings, ray-tracing methods
61 are commonly preferred both for the fluid and the foam
62 [GLR⁺06]. Although the fluid surface can be rendered
63 efficiently using ray-tracing, non-homogenous
64 phenomena such as foam require expensive volume rendering
65 techniques. In [IAAT12], the authors employed
66 a volume ray-casting method which accounts for absorption
67 and emission of radiance but neglecting light scattering
68 effects. In that method, each traced ray is sampled
69 using equally spaced intervals; and according to the
70 measured foam density at each sample point, the computed
71 radiance is attenuated. The employed ray-casting
72 approach, however, is time consuming to compute,
73 especially for scenes with many millions of foam
74 particles. The performance of volume ray-casting
75 can be significantly improved by using the GPU-based
76 method explained in [FAW10].

77 In [vdLGS09, BSW10], alternative to generating new
78 particles, selected fluid particles are visualized as foam
79 particles using GPU-based techniques for real-time
80 applications. In [BSW10], Weber number thresholding
81 is used to separate fluid and foam. Furthermore, the
82 method also takes volumetric effects into account by
83 rendering foam and fluid layers from back to front
84 order. Therefore, it can visualize effects such as foam
85 inside the fluid. Furthermore, based on the thickness
86 of the foam, it generates foam color between two user
87 defined colors. The approach, however, neglects information
88 such as occlusion and irradiance from the environment
89 when rendering foam, which limits its applicability
90 to non-photorealistic real-time renderings.

91 There also exist methods for the modeling of larger
92 scale foam effects by using air bubbles (e.g. see
93 [KVG02, KLL⁺07, HLYK08, IBAT11, BDWR12]).
94 In these works, air phase is either visualized by
95 rendering spheres [KVG02, BDWR12], or by
96 reconstructing the surface of the modeled air phase
97 [KLL⁺07, HLYK08, IBAT11]. Since we are focusing
98 on large scale scenarios, where the single air bubbles
99 inside the foam are not clearly noticeable, such
100 methods are beyond the scope of our paper.

1.3 Contribution

101 We present an efficient method for large scale foam
102 rendering. In our approach, foam is rendered using a
103 novel multi-pass rendering algorithm and finally
104 composited with the pre-rendered images of the scene
105 without foam. In comparison to volume ray-casting
106 methods that compute only absorption and emission of
107 radiance (e.g. [FAW10, IAAT12]), our approach is
108 significantly faster as the foam particles are directly
109 rendered. Furthermore, when compared to [BSW10],
110 our pipeline takes the scene occlusion and lighting into
111 account and therefore produces more convincing results
112 that can be composited with realistic renderings. Results
113 show that our new pipeline generates convincing large
114 scale foam renderings (e.g. see Fig. 1) using modern
115 GPU-based rendering architectures. 116

2 SCREEN SPACE FOAM RENDERING PIPELINE

117 As more air bubble layers implies more light scattering,
118 we relate the foam thickness to the foam intensity
119 as usually done in volume ray-casting. Later, we
120 determine the regions on screen space which should
121 receive, and therefore scatter less light using ambient
122 occlusion and attenuate the foam intensity according to
123 the occlusion factor. Afterwards, we approximate per-
124 pixel foam irradiance to colorize the foam color
125 according to the environment. Finally, the generated
126 results are composited with the rest of the scene. We
127 realized our approach using a seven pass rendering
128 algorithm. The technical steps of our pipeline (illustrated
129 in Fig. 2 and 3) can be summarized as:

- *PASS #1 and #2:* Storing eye space depth images
130 of solid and fluid meshes in two textures, which are
131 used to compute occlusion of foam fragments by
132 those primitives in the later stages. 133
- *PASS #3:* Storing an eye space depth image of the
134 foam particles in a texture, which is used in different
135 parts of our pipeline. This pass also stores a search
136 radius for each foam fragment, in whose range
137 neighboring fragments are later considered for 138

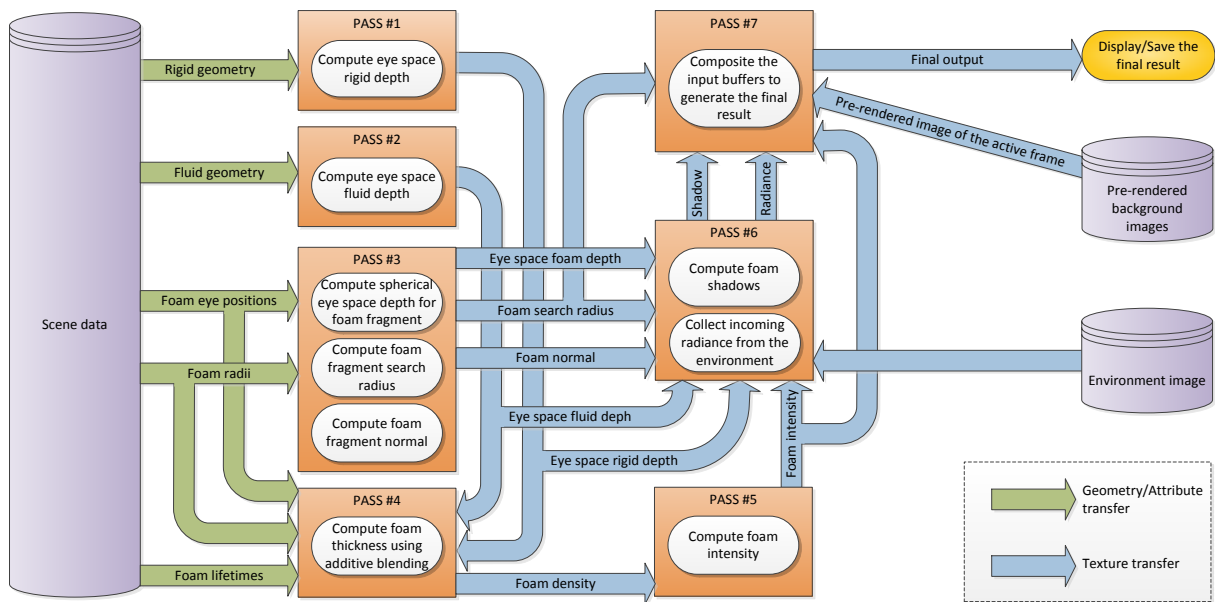


Figure 2: Diagram of our foam composition pipeline. Orange boxes denote the render passes and the arrows in between denote data flow and dependencies. For each frame, the render passes from #1 to #7 are executed. Each pass produces data explained in the enclosed rounded rectangles, which is then transferred through arrows to the subsequent passes. All of the generated textures have the same resolution as the final output.

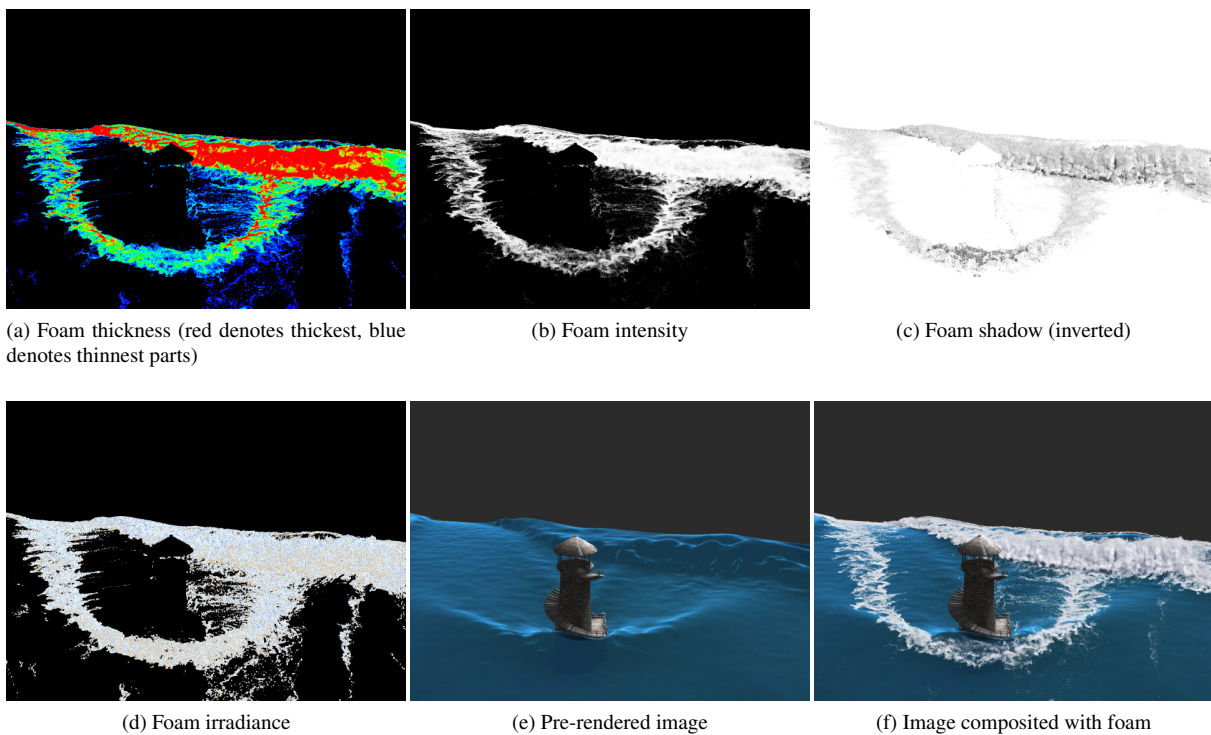


Figure 3: Some of the intermediate textures from our foam composition pipeline (a-e) and the final composited result (f).

screen space ambient occlusion and final composition (Sec. 2.1). Additionally, this pass computes a normal for each foam fragment, which is used when approximating irradiance at the fragment location.

- *PASS #4*: Accumulating foam particles via additive blending to approximate per-pixel foam thickness. This pass also discards foam fragments that are occluded by solids and attenuates foam fragments that are inside of the fluid based on the fluid transparency (Sec. 2.2).
- *PASS #5*: Conversion of per-pixel foam thickness to per-pixel foam intensity (Sec. 2.3).
- *PASS #6*: Determination of foam fragments that should receive and scatter less light using screen space ambient occlusion (SSAO) and shadow generation for such regions (Sec. 2.4.1). This pass also approximates the irradiance at each foam fragment from an environment texture if the scene is illuminated using image based lighting (Sec. 2.4.2).
- *PASS #7*: Post processing of the foam and final composition with a pre-rendered image of the scene (Sec. 2.5).

Since the first step of the pipeline is relatively straightforward, we will focus on the remaining steps throughout this section. The following render passes are implemented using OpenGL Shading Language (GLSL).

2.1 Smoothed Depth and Search Radius Computation

We use point sprites instead of spheres for rendering foam particles. A regular point sprite has the same depth values for all of its fragments. However, to produce convincing results in the later steps of our pipeline, we modify the fragment depth values similar to [vdLGS09, BSW10], such that the spherical shapes of the particles are regained.

To create the initial depth information, foam particles with ids i and radii r_i in world space are rendered with depth testing and depth masking enabled. In [IAAT12], foam particles are separated to three different types, namely: spray, surface-foam and bubble particles. For bubble particles, we use half of r_i to make them less visible. Furthermore, particle radii are randomized as $r_i = \frac{r_i}{(i \bmod 5) + 1}$ to make the particles look irregular between the scales $r_i/5$ and r_i .

The vertex shader computes eye space and projection space coordinates of the sprites and passes the resultant data to the fragment shader for further processing. In the fragment shader, the distance of the fragment position to the point sprite center is calculated using the sprite’s texture coordinates to discard fragments that are

outside of the circle. Afterwards, the flat depth values of the point sprite are transformed to spherical depth values. In this context, the first step is solving for the w coordinate of a unit sphere for the fragment’s texture coordinates in uvw space as $w = \sqrt{1 - u^2 - v^2}$, where u and v denote texture coordinates of the fragment. Subsequently, the eye space z coordinate of the fragment is simply modified as

$$\mathbf{e}_{foam_z}^{frag} = \mathbf{e}_{foam_z}^{frag} + w \cdot r_i.$$

In contrast to [vdLGS09, BSW10], we do not apply filtering to the generated depth values since it would reduce the effect of ambient occlusion.

In the same render pass, the vertex shader also projects the search radius h_i for each particle as

$$h_i = \frac{r_i}{\tan\left(\frac{\alpha}{2}\right) \left| \mathbf{e}_{foam_z}^{vert} \right|},$$

where α is the field of view of the camera and $\mathbf{e}_{foam_z}^{vert}$ denotes z coordinate of the eye position of the point sprite (i.e., distance of the sprite to the camera). Afterwards, the search radius is passed to the fragment shader as h^{frag} to be written to a texture. The depth information and the search radius are essential when rendering the SSAO pass and when doing the final composition.

This pass also computes a world space normal for each fragment \mathbf{n}_{frag} by transforming (u, v, w) using the transpose of the normal matrix, and stores the normals in a texture. Per fragment normals will be required when estimating irradiance in Sec. 2.4.2.

2.2 Thickness Estimation

Before estimating the intensity of foam at a given pixel position, we estimate the foam thickness for each pixel. In this step, foam particles are rendered again as point sprites with the spherical depth modification as in the previous render pass. Similar to [vdLGS09, BSW10], the foam fragments are blended additively to estimate thickness. Different from [vdLGS09, BSW10], however, depth buffer read and write is disabled as we do not require the frontmost particles to be visible.

As foam particles are separated to spray, surface-foam and bubble particles, we also employ this knowledge to render foam fragments differently by using a falloff function with different arguments, where the falloff is based on the fragment’s distance to the particle center in texture coordinates. The falloff function f is defined as

$$f(x, b, n, m) = \begin{cases} \left[1 - \left(\frac{x}{b}\right)^n\right]^m & \frac{x}{b} \leq 1 \\ 0 & \text{otherwise} \end{cases}, \quad (1)$$

where x is the distance to the center, b is the maximum allowed distance, and $n \geq 0$ and $m \geq 0$ are exponents which determine the shape of the function (e.g.,




Figure 4: Different forms of the falloff function given in (1) that are used in our experiments .

234 $n = 1$ and $m = 1$ result in linear falloff). When rendering spray, surface-foam and bubble fragments, we
 235 used $f_{spray} = f(x, 1, 1.5, 1)$, $f_{sfoam} = f(x, 1, 2.25, 1)$
 236 and $f_{bubble} = 1 - f(x, 1, 2, 1)$ respectively. These different falloff functions are illustrated in Fig. 4 and the
 237 corresponding intensity results are shown in Fig. 5. We
 238 preferred a larger overall intensity for surface foam particles to increase their visibility. Whereas, we preferred
 239 a comparatively smaller intensity value for the spray particles to make them relatively less visible. Further-
 240 more, we used hollow circle like structures for the bubble particles to make their appearance more convincing
 241 under water.
 242

243 In this step, the intensities of the foam particles are further modulated based on two additional factors. The
 244 first of these factors is the lifetime of the particle. For this purpose, we use $f_{lifetime} = f(l_i, 1, 2, 0.4)$, where
 245 $0 < l_i < 1$ denotes the normalized lifetime of a particle. Such a function allows a foam particle to remain
 246 visible for a sufficiently long time and fade smoothly near the end of its lifetime. Furthermore, when a particle
 247 lies in the back of the closest fluid surface (i.e. $0 < \mathbf{e}_{fluid_z}^{frag} < \mathbf{e}_{foam_z}^{frag}$, where $\mathbf{e}_{fluid_z}^{frag}$ is the eye space z co-
 248 ordinate of the fluid surface), we apply an additional falloff to its intensity, which is defined as
 249

$$f_{att} = f(\mathbf{e}_{foam_z}^{frag} - \mathbf{e}_{fluid_z}^{frag}, \eta_{max}, \eta_n, \eta_m),$$

250 with the limiting distance η_{max} , where the foam fragment completely fades to invisible, and η_n and η_m
 251 are the exponents for shaping the attenuation curve.
 252

253 At the end of this render pass, the final foam thickness values are stored in a texture (see Fig. 3a). In the next
 254 pass, the computed thickness values are processed and converted to normalized intensity values to lie between
 255 0 and 1. For all subsequent passes, a screen-filling quad is rendered to further process the relevant information
 256 that are saved in the textures.
 257

2.3 Intensity Estimation

258 As foam is composed of more bubble layers, it scatters more of the incoming light. We use this knowl-
 259 edge to relate the foam intensity proportional to foam
 260




Figure 5: Intensity distributions of different types of foam particles, namely: spray particles (left), surface foam particles (middle) and air bubble particles (right).




Figure 6: Different forms of the sigmoid function that can be applied to the accumulated foam densities. The function can be used to create different distributions as well. For instance, to reduce the intensities below some threshold, $\rho_{exp} \geq 2$, can be used. We use the $t(\rho, 3, 1.25)$ form in our experiments.

273 thickness. A texel from the previous render pass may have any value between $[0, \infty)$. In this render pass, we
 274 scale the values taken from that texture to the interval $[0, 1]$. However, scaling the values linearly to the target
 275 interval would make sparse areas invisible. We expect the foam to become completely opaque after some
 276 thickness threshold. Therefore, to increase the effective range of the thinner regions, to reduce the range
 277 of thicker regions and to normalize the intensities, we define the following sigmoid function t to non-linearly
 278 scale a pixel thickness value ρ as
 279

$$t(\rho, \rho_{mod}, \rho_{exp}) = \frac{\rho^{\rho_{exp}}}{\rho_{mod} + \rho^{\rho_{exp}}},$$

280 where $\rho_{mod} > 0$ and $\rho_{exp} > 0$ control how fast the function grows. Note that if $\rho > 0$ and $\rho_{exp} > 0$, $0 < t < 1$.
 281 t is illustrated in Fig. 6 for different parameters. Furthermore, Fig. 7-top shows the effect of using different
 282 ρ_{mod} values.
 283

284 At the end of this step, the normalized intensities are saved in a texture, which will be used in the following
 285 steps (see Fig. 3b).
 286

2.4 Foam Radiance Estimation

287 Since foam is composed of many transparent layers of air bubbles, light can travel through it and then scatter.
 288 Until the current stage of our pipeline, we assume that foam scatters light uniformly, where the intensity
 289

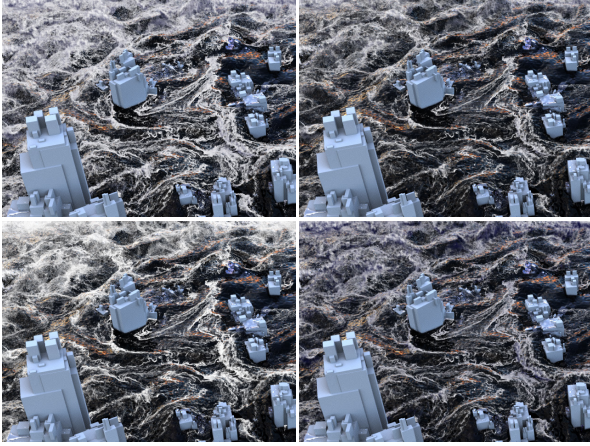


Figure 7: Application of different parameters for the setting presented in Fig. 1-middle. top-left: $\rho_{mod} = 1$; top-right: $\rho_{mod} = 5$; bottom-left: $AO_{ShScale} = 0.1$; bottom-right: $AO_{ShScale} = 2$.

297 of the light was only related to the foam thickness. In
 298 this section, we determine the regions which should receive, and therefore scatter less light using ambient occlusion (AO), and generate shadows for these regions
 299 (Sec. 2.4.1). Furthermore, the intensities that are computed in the previous section do not employ any knowledge about the actual illumination that comes from the scene. In this render pass, we will also use a very rough screen space approximation of the irradiance from the surrounding environment, which is used to colorize the foam fragments (Sec. 2.4.2).
 300
 301
 302
 303
 304
 305
 306
 307

308 This render pass again gets the textures that have been
 309 computed in the previous step as input and computes two additional textures, one for the shadow and another for the illumination of the foam (see Fig. 3).
 310
 311

312 2.4.1 Shadow Generation

313 As object space AO methods (e.g. [ZIK98, Bun05, RWS⁺06]) are very expensive to compute, especially for complex dynamical phenomena such as foam, we investigated SSAO techniques [TCM06, Mit07, SA07, RGS09, BS09, HL10]. Finally, we decided to build our SSAO approach upon the basic concept explained in [Mit07] because of its efficiency and simplicity. One important difference of our method in comparison to [Mit07] is that we apply multiple sample collection iterations to capture both small scale and large scale occlusions. Instead of increasing search radii, [HL10] used multiple depth maps with decreasing resolution to achieve the same effect. The search radii and total number of passes are controlled by three parameters: the initial search radius factor $AO_{InitSRFac}$, which is a factor for h^{frag} to capture small scale occlusions; the search radius increment factor $AO_{SRIncFac}$, which is another factor for h^{frag} to determine how much the search radius increases in each sample collection step; and finally $AO_{#Passes}$, which limits the total number of SSAO
 314
 315
 316
 317
 318
 319
 320
 321
 322
 323
 324
 325
 326
 327
 328
 329
 330
 331
 332

333 passes. For each fragment, 3d samples are generated
 334 within the fragment search radius:

$$h_{pass}^{frag} = h^{frag} (AO_{InitSRFac} + AO_{SRIncFac} \cdot AO_{Pass}),$$

335 where AO_{pass} increases by 1 in each sample collection
 336 pass and $AO_{pass} \leq AO_{#Passes}$. In our experiments we
 337 used: $AO_{InitSRFac} = 1$, $AO_{SRIncFac} = 7$ and $AO_{#Passes} =$
 338 3.

339 The total number of samples v in each sample collec-
 340 tion pass is controlled by a user defined sampling den-
 341 sity parameter AO_{SDens} as

$$v = \text{clamp} \left(\frac{3}{4} \pi h_{pass}^{screen^3} AO_{SDens}, AO_{\#MinSamp}, AO_{\#MaxSamp} \right),$$

342 where h_{pass}^{screen} is the search radius projected to fragment
 343 coordinates. Since h^{frag} can be very small for distant
 344 fragments, a minimum value $AO_{\#MinSamp}$ is used for v .
 345 An upper limit $AO_{\#MaxSamp}$ is also introduced to pre-
 346 vent too many samples from being generated for frag-
 347 ments that are very close to the viewer. In our experi-
 348 ments, we used $AO_{SDensity} = 0.5$, $AO_{\#MinSamp} = 16$ and
 349 $AO_{\#MaxSamp} = 512$. The samples are created inside a
 350 cube in the range $[-1, 1]$ on all axes using the Hal-
 351 ton sampling algorithm with a constant seed [Hal64],
 352 which produces low-discrepancy sequences. Subse-
 353 quently, the samples are mapped to a sphere by simply
 354 neglecting the samples that lie outside of the sphere in
 355 the range $[-1, 1]$.

356 Additionally, the occlusion contribution λ of a sample s
 357 depends on its distance to the fragment and we compute
 358 it using a quadratic falloff as

$$\lambda = (1 - |s|)^2.$$

359 Furthermore, if a sample is occluded by a fragment with
 360 a distance larger than the user defined $AO_{MaxOcclDist}$,
 361 the sample does not contribute to the visibility of the
 362 fragment. This effect is necessary to prevent occlusion
 363 by distant fragments and is controlled using a quadratic
 364 falloff function as

$$\delta = \max \left[\left(1 - \frac{|e_{foam_z}^{frag} - s_z|}{AO_{MaxOcclDist}} \right), 0 \right]^2,$$

365 where $AO_{MaxOcclDist} = 5$ is used in our experiments.
 366 The sample s is used to look up the occlusion in eye
 367 space by other fragments (e.g. foam, fluid and solid
 368 fragments) in the scene. Based on the knowledge col-
 369 lected so far, the occlusion k of a sample is defined as

$$k = \begin{cases} 1 & \left[(s_z > e_{foam_z}^{frag} \vee s_z > e_{fluid_z}^{frag} \vee s_z > e_{rigid_z}^{frag}) \wedge (0 < \delta < 1) \right] \\ 0 & \text{otherwise} \end{cases},$$

370 which basically states that; if a sample is occluded by
 371 any other fragment in the scene and if the occlusion dis-
 372 tance is not larger than $AO_{MaxOcclDist}$, the sample is oc-
 373 cluded.

374 Afterwards, we compute the occlusion factor ω of a
 375 fragment as

$$\omega = \frac{\sum_{AO_{pass}=1}^{AO_{\#Passes}} (\sum_{i=1}^V \lambda_i \cdot \delta_i \cdot k_i \cdot a_i)}{\sum_{AO_{pass}=1}^{AO_{\#Passes}} (\sum_{i=1}^V \lambda_i)},$$

376 where for the pass AO_{pass} , i iterates over all generated
 377 samples which are inside the render area, and a_i is the
 378 transparency of the sampled fragment, which is equiv-
 379 alent to t_i for foam fragments. For rigid and fluid frag-
 380 ments, a_i is equivalent to the fragment’s transparency.
 381 Additionally, if there are multiple overlapping transpar-
 382 ent fragments at a sample position, a_i is computed by
 383 adding all of the transparency values.

384 Finally, so as to be more flexible about the appearance
 385 of the generated shadows, we compute the final shadow
 386 value ζ clamped into $[0, 1]$ as

$$\zeta = \text{clamp} \left[(\omega \cdot AO_{ShScale})^{AO_{ShExp}} + AO_{ShOffset}, 0, 1 \right]$$

387 which is controlled by three self-explanatory user de-
 388 fined parameters. In the presented scenarios, we used:
 389 $AO_{ShScale} = 1$, $AO_{ShExp} = 1.5$ and $AO_{ShOffset} = -0.05$.
 390 The ambient occlusion step especially improves the re-
 391 gions that have similar intensities, which would look
 392 totally flat otherwise (e.g., see Fig. 8, top-middle). Fur-
 393 thermore, Fig. 7-bottom shows the effect of different
 394 $AO_{ShScale}$ values. The computed ζ values are written
 395 to a texture to be further used by the final composition
 396 step (see Fig. 3c).

397 2.4.2 Irradiance

398 If the scene is illuminated using image based lighting,
 399 we approximate the direct illumination of each foam
 400 fragment by looking up the environment map that has
 401 been used as the light source. Using the fragment nor-
 402 mal \mathbf{n}^{frag} , we create a hemisphere around the normal
 403 and use the already generated samples from the SSAO
 404 step to create direction vectors \mathbf{n}_i^{sample} that are used for
 405 looking up the intensity $\mathbf{P} = (r, g, b)$ at an environment
 406 map position. Finally, the irradiance that is coming
 407 from the environment to a fragment location is simply
 408 computed in a cosine weighted fashion as

$$\mathbf{I} = \left(\frac{\sum_{i=1}^V \mathbf{P} \cdot (\mathbf{n}_i^{sample} \cdot \mathbf{n}^{frag})}{\sum_{i=1}^V i} \right),$$

409 where i iterates only over the samples that are generated
 410 for the first sample collection pass. The sole purpose of
 411 this step is to reflect the hue of the environment onto

412 the foam fragments to make the foam not look too dis-
 413 tinct from the rest of the scene. Finally, the computed
 414 \mathbf{I} values are written to another texture to be used by
 415 the next and the final render pass (see Fig. 3d). The
 416 performance of this step can be improved by using an
 417 irradiance environment map and making color look-up
 418 once for every \mathbf{n}^{frag} .

419 2.5 Composition

420 In this render pass, the information that has been cre-
 421 ated in the previous steps and the pre-rendered images
 422 of the scene without foam are composited to generate a
 423 final image of the scene with foam (see Fig. 2).

424 Depending on the user defined $AO_{\#MaxSamples}$, the
 425 shadow and radiance values computed in the previous
 426 section can include high frequency noise. In order
 427 to alleviate this problem, we apply Gaussian blur
 428 with a filter radius of $\frac{3}{2}h_{pass}^{screen}$ to both textures to
 429 generate per-pixel $\zeta_{filtered}$ and $\mathbf{I}_{filtered}$ before doing the
 430 composition.

431 Afterwards, to compute a final shadow color ζ_{final} for
 432 a pixel, the filtered shadow values are modulated with a
 433 user defined color $\mathbf{C}_{ShadowColor}$ and clamped into $[0, 1]$
 434 as

$$\zeta_{final} = \text{clamp} \left[(\mathbf{c}_{white} - \mathbf{C}_{ShadowColor}) \odot \zeta_{filtered}, 0, 1 \right],$$

435 where $\mathbf{c}_{white} = (1, 1, 1)$, and \odot denotes component-
 436 wise multiplication. We select $\mathbf{C}_{ShadowColor}$ similar to
 437 the visible color of the fluid that the foam is generated
 438 on, and it was chosen in our experiments as $(0, 0, 0.2)$
 439 because of the dark blue appearance of the fluids in our
 440 renderings. Since ζ_{final} will be subtracted when do-
 441 ing the composition, the $\mathbf{C}_{ShadowColor}$ term is subtracted
 442 from white to invert it. From our experiences, coloriz-
 443 ing shadows makes the foam blend better with the un-
 444 derlying fluid.

445 As foam is composed of many air-liquid interfaces, it
 446 has a very large scattering albedo which causes it to
 447 scatter most of the incoming light, but absorb only a
 448 small amount of it. Therefore, it is usually observed
 449 very bright. We control this phenomenon by linearly
 450 scaling the irradiance values \mathbf{I} using a user defined pa-
 451 rameter $C_{IrrScale}$, whose value depends on the desired
 452 foam brightness and the color range of the environment
 453 map used. Afterwards, we clamp the resulting color
 454 into the $[0, 1]$ interval to compute

$$\mathbf{I}_{final} = \text{clamp} (C_{IrrScale} \cdot \mathbf{I}_{filtered}, 0, 1).$$

455 Finally, the composited pixel color c is computed as

$$\mathbf{c} = (1 - \iota) \mathbf{c}_{bg} + \iota (\mathbf{I}_{final} - \zeta_{final})$$

456 where \mathbf{c}_{bg} is the color at the corresponding pixel po-
 457 sition of the background image on which the foam is
 458 composited (see Fig. 3f).




Figure 8: Comparisons of our method (top) to volume ray-casting that computes emission and absorption only (bottom). As our method approximates shadows in concave regions, the foam looks more volumetric and detailed. The scenes are named from left to right as: Wave, Lighthouse and Ship.

	# Foam particles	Resolution	Average foam rendering time per frame			
			Ray-casting [IAAT12]	Ray-casting [FAW10]	Ours (intensity only)	Ours (total)
Wave	up to 820K	800 × 600	2 min 10 s	235 ms	8 ms	52 ms
Ship	up to 9M	800 × 600	4 min 20 s	760 ms	16 ms	102 ms
Lighthouse	up to 15M	800 × 600	7 min 3 s	1 s	21 ms	150 ms
Flood	up to 29M	1280 × 960	16 min 19 s	1.7 s	33 ms	235 ms

Table 1: Performance analysis of the example scenes.

3 RESULTS

459 In this section, we demonstrate the versatility of our
 460 approach in different animation sequences. For all
 461 presented scenes, the underlying fluid has been simu-
 462 lated using the methods referred in [IAAT12], and
 463 the fluid surfaces have been reconstructed based on
 464 [SSP07, AIAT12, AAIT12]. The scenes were rendered
 465 using mental ray [NV11] on an Intel Xeon X5690 CPU
 466 with 12 GB RAM, and the foam composition pipeline
 467 was implemented using GLSL and ran on an NVIDIA
 468 480 GTX GPU with 1.5 GB RAM. The ray-casting
 469 code used in [IAAT12] was implemented as a mental
 470 ray shader and ran on the CPU, and an optimized ver-
 471 sion based on [FAW10] was implemented on the GPU.
 472 All scenes were illuminated using image based lighting
 473 with a clear sky environment map.

474 For all scenes, foam was simulated using [IAAT12] and
 475 the same foam data were used for the rendering com-
 476 parisons to [IAAT12]. For the comparisons shown in
 477 Fig. 8, the ray-casting technique explained in [IAAT12]
 478 took 9 s to 20 min depending on the complexity of the
 479 frame, excluding the other scene geometry and load-
 480 ing of the foam data. Using the optimized volume ray-
 481 casting scheme, the computation time has been reduced
 482 down to 90 ms to 2.5 s. Using our pipeline, the foam

483 rendering of a frame took 30 ms to 270 ms depending
 484 on the complexity of the foam in the scene being ren-
 485 dered, excluding the time spent for loading of the foam
 486 data from secondary storage to the GPU memory. The
 487 results produced by using a basic volume ray-casting
 488 scheme that only accounts for absorption and emission
 489 of radiance is similar to the results we achieve exclu-
 490 ding the additional effects that are described in Sec. 2.4
 491 (see also Fig. 3b). Excluding those additional effects,
 492 our pipeline took between 5 ms to 39 ms per frame. See
 493 Table 1 for additional information about each scene. As
 494 our pipeline also takes additional effects into account
 495 (i.e. ambient occlusion and irradiance estimation), our
 496 presented foam renderings look volumetric and blend
 497 with the rest of the scene (see Fig. 8). Note that in
 498 [IAAT12], the fluid surface has been constructed only
 499 for the fluid particles that have more than five neigh-
 500 bors. For our comparisons to [IAAT12], however, we
 501 used the whole fluid surface for our renderings to bet-
 502 ter estimate the SSAO of the foam by the fluid surface.
 503 Therefore, differences between the two fluid surfaces
 504 can be noticeable.

505 For all of our scenes, most of the rendering time has
 506 been spent on the foam radiance estimation pass (be-

507 tween 50-80%). Whereas, the computational overheads
508 of the rest of the render passes were significantly lower.

4 DISCUSSION AND FUTURE WORK

509 Taking a closer look at sea foam from a distance less
510 than a few meters, one may observe the underlying air
511 bubbles at varying sizes which form the foam. Render-
512 ing of such scenarios is not handled by our approach.
513 However, using an air bubble generation technique like
514 [BDWR12] for such close-ups might be an interesting
515 direction for future research.

516 For scenes where most of the light is coming from a
517 specific direction at shallow angles (e.g. sunset scenar-
518 ios), the currently employed SSAO based shadow gen-
519 eration technique can fail to capture the resultant po-
520 tentially large shadows cast by distant objects. For such
521 cases, an explicit shadow generation algorithm which
522 can handle image based lighting such as the one ex-
523 plained in [CK09], or explicit shadow source selection
524 as discussed in [Bjo04] can be employed. Since we as-
525 sume that foam scatters most of the incident light ran-
526 domly, we omitted Fresnel effect. However, it might
527 be desirable to make the foam reflect the environment,
528 when it is viewed from a shallow angle.

529 Our algorithm neglects many physical effects that could
530 be otherwise simulated by using modern ray-tracing
531 techniques. Those effects include; scattering of light
532 inside the foam, influence of the foam on the appear-
533 ance of the surrounding objects and vice versa. How-
534 ever, for large scale scenarios (e.g. as in Fig 1), those
535 effects have less significance on the appearance of the
536 foam, and our approximations can efficiently gener-
537 ate convincing results. However, for close-ups, the
538 effects that we have omitted have more significance
539 on the final outcome. For those cases, using a vol-
540 ume ray-casting method that simulates light scatter-
541 ing can definitely yield more convincing results (e.g.
542 [RNGF03, GLR⁺06]).

543 Although we demonstrated our rendering scheme only
544 for the particle data generated by the method explained
545 in [IAAT12], we believe that our pipeline is mostly ap-
546 plicable to the rendering of other particle based foam
547 simulation techniques.

5 CONCLUSION

548 We presented an efficient, screen-space foam rendering
549 pipeline that can render large particle-based foam data
550 sets on the GPU. Our approach uses a multi-pass ren-
551 dering scheme, where different effects are added to the
552 foam rendering incrementally, and the final foam ren-
553 dering is composited with a pre-rendered image of the
554 scene. The presented method can be used as an efficient
555 alternative to ray-casting techniques for the rendering
556 of large scale particle-based foam data.

6 REFERENCES

- [AAIT12] G. Akinci, N. Akinci, M. Ihmsen, and M. Teschner. An efficient surface reconstruction pipeline for particle-based fluids. In *Workshop on Virtual Reality Interaction and Physical Simulation*, pages 61–68. The Eurographics Association, 2012.
- [AIAT12] Gizem Akinci, Markus Ihmsen, Nadir Akinci, and Matthias Teschner. Parallel surface reconstruction for particle-based fluids. *Computer Graphics Forum*, 31:1797–1809, 2012.
- [BDWR12] O. Busaryev, T.K. Dey, H. Wang, and Z. Ren. Animating bubble interactions in a liquid foam. *ACM Transactions on Graphics (TOG)*, 31(4):63, 2012.
- [Bjo04] K. Björke. Image-based lighting, 2004.
- [BS09] L. Bavoil and M. Sainz. Multi-layer dual-resolution screen-space ambient occlusion. In *SIGGRAPH 2009: Talks*, page 45. ACM, 2009.
- [BSK⁺07] R. Bredow, D. Schaub, D. Kramer, M. Hausman, D. Dimian, and R.S. Duguid. Surf’s up: the making of an animated documentary. In *International Conference on Computer Graphics and Interactive Techniques: ACM SIGGRAPH 2007 courses: San Diego, California*, volume 5, 2007.
- [BSW10] F. Bagar, D. Scherzer, and M. Wimmer. A layered particle-based fluid model for real-time rendering of water. In *Computer Graphics Forum*, volume 29, pages 1383–1389. Wiley Online Library, 2010.
- [Bun05] M. Bunnell. Dynamic ambient occlusion and indirect lighting. *Gpu gems*, 2(2):223–233, 2005.
- [CK09] Mark Colbert and Jaroslav Krivanek. Real-time dynamic shadows for image-based lighting. *ShaderX 7 - Advanced Rendering Techniques*, page Section 4.3, 2009.
- [CM10] N. Chentanez and M. Müller. Real-time simulation of large bodies of water with small scale details. In *Proc. of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 197–206, 2010.
- [FAW10] R. Fraedrich, S. Auer, and R. Westermann. Efficient high-quality volume rendering of sph data. *Visualization and Computer Graphics, IEEE Transactions on*, 16(6):1533–1540, 2010.
- [GLR⁺06] W. Geiger, M. Leo, N. Rasmussen, F. Losasso, and R. Fedkiw. So real it’ll make you wet. In *ACM SIGGRAPH*, 2006.
- [Hal64] J.H. Halton. Algorithm 247: Radical-inverse quasi-random point sequence. *Communications of the ACM*, 7(12):701–702, 1964.
- [HL10] T.D. Hoang and K.L. Low. Multi-resolution screen-space ambient occlusion. In *Proceedings of the 17th ACM Symposium on Virtual Reality Software and Technology*, pages 101–102. ACM, 2010.
- [HLYK08] Jeong-Mo Hong, Ho-Young Lee, Jong-Chul Yoon, and Chang-Hun Kim. Bubbles alive. *ACM Trans. Graph.*, 27(3):48:1–48:4, August 2008.
- [hyb11] Next Limit Technologies: Realfow 2012, Hybrid. White Paper. 2011.
- [IAAT12] Markus Ihmsen, Nadir Akinci, Gizem Akinci,

- and Matthias Teschner. Unified spray, foam and air bubbles for particle-based fluids. *The Visual Computer*, pages 1–9, 2012. 10.1007/s00371-012-0697-9.
- [IBAT11] M. Ihmsen, J. Bader, G. Akinci, and M. Teschner. Animation of air bubbles with sph. In *International Conference on Graphics Theory and Application*, pages 225–234, 2011.
- [KLL⁺07] Byungmoon Kim, Yingjie Liu, Ignacio Llamas, Xiangmin Jiao, and Jarek Rossignac. Simulation of bubbles in foam with the volume control method. *ACM Trans. Graph.*, 26(3), July 2007.
- [KVG02] Hendrik Kück, Christian Vogelgsang, and Günther Greiner. Simulation and rendering of liquid foams. In *In Proc. Graphics Interface '02*, pages 81–88, 2002.
- [LTKF08] F. Losasso, J. Talton, N. Kwatra, and R. Fedkiw. Two-Way Coupled SPH and Particle Level Set Fluid Simulation. *IEEE Transactions on Visualization and Computer Graphics*, 14(4):797–804, 2008.
- [Mit07] M. Mittring. Finding next gen: Cryengine 2. In *ACM SIGGRAPH 2007 courses*, pages 97–121. ACM, 2007.
- [MMS09] V. Mihalef, D. Metaxas, and M. Sussman. Simulation of two-phase flow with sub-scale droplet and bubble effects. In *Computer Graphics Forum*, volume 28, pages 229–238. Wiley Online Library, 2009.
- [MSD07] M. Müller, S. Schirm, and S. Duthaler. Screen Space Meshes. In *Proceedings of ACM SIGGRAPH / EUROGRAPHICS Symposium on Computer Animation (SCA)*, 2007.
- [NV11] NVIDIA ARC. mental ray 3.9 [software]. <http://www.mentalimages.com/products/mental-ray/about-mental-ray.html>, 2011.
- [RGS09] T. Ritschel, T. Grosch, and H.P. Seidel. Approximating dynamic global illumination in image space. In *Proceedings of the 2009 symposium on Interactive 3D graphics and games*, pages 75–82. ACM, 2009.
- [RNGF03] N. Rasmussen, D.Q. Nguyen, W. Geiger, and R. Fedkiw. Smoke simulation for large scale phenomena. *ACM Transactions on Graphics (TOG)*, 22(3):703–707, 2003.
- [RWS⁺06] Z. Ren, R. Wang, J. Snyder, K. Zhou, X. Liu, B. Sun, P.P. Sloan, H. Bao, Q. Peng, and B. Guo. Real-time soft shadows in dynamic scenes using spherical harmonic exponentiation. In *ACM Transactions on Graphics (TOG)*, volume 25, pages 977–986. ACM, 2006.
- [SA07] P. Shanmugam and O. Arikan. Hardware accelerated ambient occlusion techniques on gpus. In *Proceedings of the 2007 symposium on Interactive 3D graphics and games*, pages 73–80. ACM, 2007.
- [SSP07] B. Solenthaler, J. Schläfli, and R. Pajarola. A unified particle model for fluid-solid interactions. *Computer Animation and Virtual Worlds*, 18(1):69–82, 2007.
- [TCM06] M. Tarini, P. Cignoni, and C. Montani. Ambient occlusion and edge cueing for enhancing real time molecular visualization. *Visualization and Computer Graphics, IEEE Transactions on*, 12(5):1237–1244, 2006.
- [TFK⁺03] T. Takahashi, H. Fujii, A. Kunimatsu, K. Hiwada, T. Saito, K. Tanaka, and H. Ueki. Realistic Animation of Fluid with Splash and Foam. *Computer Graphics Forum*, 22(3):391–400, 2003.
- [vdLGS09] 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.
- [YT10] J. Yu and G. Turk. Reconstructing surfaces of particle-based fluids using anisotropic kernels. In *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pages 217–225. Eurographics Association, 2010.
- [ZB05] Y. Zhu and R. Bridson. Animating sand as a fluid. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*, pages 965–972, New York, NY, USA, 2005. ACM Press.
- [ZIK98] S. Zhukov, A. Iones, and G. Kronin. An ambient light illumination model. *Rendering techniques*, 98:45–55, 1998.