

# Real-time Rendering of Enhanced Shallow Water Fluid Simulations

Jesús Ojeda<sup>a</sup>, Antonio Susín<sup>b</sup>

<sup>a</sup>Dept. LSI, Universitat Politècnica de Catalunya

<sup>b</sup>Dept. MA1, Universitat Politècnica de Catalunya

---

## Abstract

The visualization of simulated fluids is critical to understand their motion, with certain light effects restricted or with added computational complexity in the implementation if real-time simulation is required. We propose some techniques that improve the rendering quality of an enhanced shallow waters simulation. To improve the overall appeal of the fluid representation, lower scale details are added to the fluid, coupling external non-physical simulations, and advecting generated surface foam. We simulate caustics by raytracing photons in light and screen-space, and apply refraction and reflections also in screen-space, through a number of render passes. Finally, it is shown how a reasonably sized fluid simulation is executed and rendered at interactive framerates with consumer-level hardware.

*Keywords:* real-time reflections and refractions, real-time caustics, fluid rendering

---

## 1. Introduction

Photorealistic rendering is still quite demanding for interactive applications due to its computational complexity. Otherwise, given enough time, offline renderers can easily generate this kind of imagery, usually using some algorithm of the ray-tracing family.

In the real-time field, however, GPUs are used which implement rasterization algorithms. These algorithms rely on high coherency for the operations executed, which impose some constraints to simulate light as a raytracer could do, by simulating each separate light beam. For this reason, photorealistic rendering is achieved at interactive framerates by simplifying the algorithms used or even using tricks that are perceptually feasible.

This simulation of light behaviour is required if a realistic fluid visualization is pursued. Liquids, in their vast majority, exhibit reflection and refraction effects, which in turn may also result in caustics. Assuming a visualization over the fluid, for great volumes of water, as open sea scenes, the major visual effects one could expect may be the refraction of the underlying terrain with projected caustics, as well as reflected scenery from above the fluid.

With present fluid simulations being performed in GPUs at interactive framerates, we also need realistic visualizations which reproduce these effects of the light. In our case, we start from a heightfield fluid simulation enhanced with particles for the simulation of splashes in breaking wave conditions, like the ones proposed in [1] or [2], which are also fully coupled with dynamic objects. From there, we aim to provide these expected, light-based effects, namely refractions, reflections and caustics.

As the fluid simulation mesh may have lower resolution than that expected for high quality results, it is also improved with other techniques as lower scale details and surface foam advection which effectively increase the general appeal of the rendered scenes. The key contributions we propose are

- An extension of the technique from [3] to screen-space, following an initial photon search in light-space, as well as some other modifications for the simulation of caustics.
- A screen-space technique to simulate refractions and reflections, based in raycasting through depth-maps.
- Texture-based techniques for additional surface effects using FFT ocean simulation or Perlin noise, as well as the advection of surface foam generated at the splash particles reintroduction.

The result of these contributions is exemplified in Figure 1, and how they are interlaced as an overall algorithm can be seen in Figure 2.

### 1.1. Related Work

There are two common approaches to simulate fluids: eulerian and laplacian. The first simulate the fluid inside a grid. In the second, the fluid is implicitly represented by a particle system. For a full 3D fluid simulation we can find many references of both approaches but for brevity reasons we refer the reader to [4], [5] and references therein for greater fluid overviews.

In the specific case of eulerian fluid simulation, the fluid is usually represented as a scalar field and its visualization is done by raycasting the volume or by using mesh-extracting techniques like marching cubes for further use. Nevertheless, 3D full simulation can be still quite costly, so other solutions as

---

*Email addresses:* [jojeda@lsi.upc.edu](mailto:jojeda@lsi.upc.edu) (Jesús Ojeda),  
[toni.susin@upc.edu](mailto:toni.susin@upc.edu) (Antonio Susín)




Figure 1: Caustics on the underlying terrain can be seen through the refractive surface of the fluid.

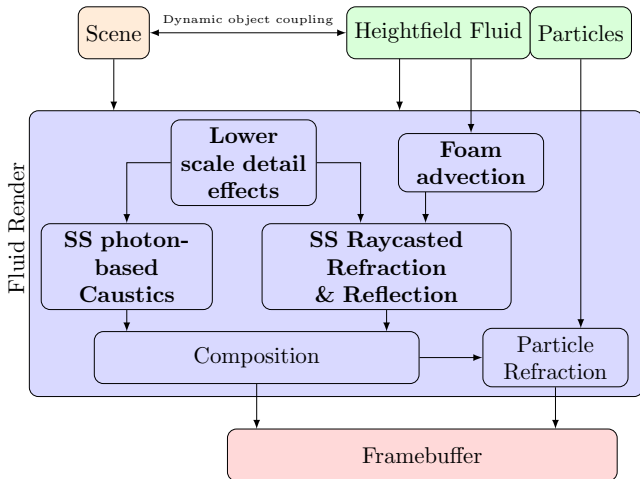


Figure 2: Pipeline of the different parts involved in the rendering of enhanced shallow water simulation with particles, providing our screen space photon-based caustics, screen-space raycasted refraction and reflection, as well as other surface effects as lower scale details and foam advection.

61 heightfield representations are more commonly used in the  
 62 industry of real-time applications. Such simulations can come  
 63 from procedural methods like the FFT ocean simulation [6],  
 64 wave trains [7] or even physical frameworks as the Shallow Wa-  
 65 ter equations [8, 2, 1]. Their result, can be easily represented as  
 66 a triangle mesh, where vertex heights are provided by the own  
 67 simulation.

68 As these grid-based approaches have a fixed resolution, in  
 69 order to increase the perceived level of detail, other techniques  
 70 have been applied as the advection of additional textures to sim-  
 71 ulate flow [9], coupled with normal mapping as in [2].

72 To finally visualize the fluid, several light-induced effects  
 73 have to be considered. One of these effects are caustics, which  
 74 are a very distinguishable effect from any refractive or reflective  
 75 surface. Starting from Kajiya’s work [10], caustics have been  
 76 traditionally implemented with global illumination techniques  
 77 like pathtracing, the metropolis light transport method [11] or  
 78 photon mapping [12]. These techniques require a high count of  
 79 rays or photons to achieve soft caustics, which relegate them to  
 80 off-line rendering, although there are already GPU implemen-  
 81 tations of some of them like, e.g., [13, 14].

82 In the real-time domain, [15] was the first to explore caus-  
 83 tics using synthetic texture maps; although inaccurate, they were  
 84 visually compelling. Nevertheless, to achieve physically real-  
 85 istic results, the more recent techniques are inspired in path-  
 86 tracing methods and can be generally classified in two groups.  
 87 In the first group, techniques like, e.g. [16, 17, 3, 18], render  
 88 from light and create caustic maps, similar to photon maps, but  
 89 used like shadow maps; reprojected in camera space in order to  
 90 lit the visible pixels that receive caustics. In the second group,  
 91 the caustics are traced back from the receiver object to the light  
 92 through limited areas on the refractive surface as in [19, 20],  
 93 which usually require the receiver to be planar as a simplifica-  
 94 tion.

95 Similarly, for the simulation of refractive or reflective ma-  
 96 terials, the ground truth may be reached with the usual path-  
 97 tracing techniques but in the real-time domain trick techniques,  
 98 like [21] which apply a random offset to the refracted vector,  
 99 are commonly used. [16] on the other hand relies on environ-  
 100 ment maps as distance impostors to achieve approximate refraction.  
 101 However, for a more physically accurate approach, the  
 102 more complete techniques involve tracing rays through depth  
 103 maps. In this sense, [22] simulated refraction using front and  
 104 back depth maps, while later [23] improved on the previous  
 105 technique by repeating the search in between depth maps to  
 106 simulate total internal refraction. [24] also worked upon [22],  
 107 improving it with depth corrections, impostors and caustics.

108 In contrast, we provide a full system for caustics, reflections  
 109 and refractions as well as other effects to complete the fluid ren-  
 110 dering. For the caustics, we improve upon [3], adding a second  
 111 raycast phase in screen space (from the camera) to the first one  
 112 used in light space. Furthermore, we simplify their approach  
 113 by not generating a caustics map, but splatting the photons on  
 114 the receiving geometry. In the reflection and refraction case,  
 115 we specialize the refraction approach of [22, 23] to our fluid  
 116 scenes: we render the geometry over and below the fluid sep-  
 117 arately and apply the same raycast algorithm to both buffers,

118 combining the results using Fresnel terms.

## 119 2. Fluid simulation

In our case, we use the fluid solver from [1], the simulation algorithm is based on the the Lattice Boltzmann Method (LBM) for Shallow Waters. The fluid is simulated in a grid and the interactions between the fluid molecules (described as distribution functions  $f_i$ ) are modeled as collisions. Using the D2Q9 model, the LBM with the popular BGK collision operator [25] can be defined with the following equation:

$$f_i(\mathbf{x} + \mathbf{e}_i \Delta t, t + \Delta t) = f_i(\mathbf{x}, t) - \omega(f_i - f_i^{eq}) + \mathcal{F}_i, \quad (1)$$

where  $\omega$  is the relaxation parameter related to the viscosity of the fluid,  $\mathcal{F}_i$  are external forces and  $f_i^{eq}$  is the equilibrium distribution function defined as

$$f_i^{eq}(h, \mathbf{u}) = \begin{cases} h \left(1 - \frac{5}{6}gh - \frac{2}{3}\mathbf{u}^2\right), & i = 0, \\ \lambda_i h \left(\frac{gh}{6} + \frac{\mathbf{e}_i \cdot \mathbf{u}}{3} + \frac{(\mathbf{e}_i \cdot \mathbf{u})^2}{2} - \frac{\mathbf{u}^2}{6}\right), & i \neq 0, \end{cases} \quad (2)$$

120 where  $\lambda_i = 1$  for  $i = 1..4$  and  $\lambda_i = 1/4$  for  $i = 5..8$ .  $g$  is  
121 the gravity and  $h$  and  $\mathbf{u}$  are the fluid properties: height level  
122 from the underlying terrain and velocity, respectively. They are  
123 calculated as


$$h(\mathbf{x}, t) = \sum_i f_i, \quad (3)$$

$$\mathbf{u}(\mathbf{x}, t) = \frac{1}{h} \sum_i \mathbf{e}_i f_i. \quad (4)$$

124 This basic model is enhanced by applying breaking wave  
125 conditions, an additional particle system and two-way object  
126 coupling similarly to [2]. Except the coupling with external ob-  
127 jects, simulated with the Bullet Physics library, the whole simu-  
128 lation is executed in CUDA and achieves interactive timerates.  
129 We refer the reader to [1] for a full review on the simulated fluid  
130 system, the breaking wave example shown in Figure 3.

131 As the fluid is provided as a heightfield, its basic visual-  
132 ization can be a triangle mesh representing the whole domain,  
133 being the vertices equally displaced in the  $xz$  plane and their  
134  $y$  coordinate the value of the heightfield at that point. We use  
135 this representation, and apply some techniques that enable more  
136 complex visual effects as caustics and refraction, which aren't  
137 restricted to this Shallow Waters simulation and may be applied  
138 to other refractive/reflective surfaces. These techniques are ex-  
139 plained in the next sections.

140 For the particles, we render them as points, expanded to  
141 quadrilaterals and use depth and normal replacement, similarly  
142 to [26]. For refraction, methods like [21, 27] can be used. In  
143 our case we use the first one, where an arbitrary offset is ap-  
144 plied to the refracted vector from the particle surface normal  
145 and used to look up at the framebuffer; although the results of  
146 this arbitrary offset on the refracted vectors are not physically  
147 correct, they are perceptually feasible and simpler to implement  
148 than the latter one, for example.



(a) FFT ocean simulation. (b) Noise, 4 octaves of a fractal sum.

Figure 4: Adding lower scale details to the fluid surface by normal mapping. Refraction and reflection are deactivated.

## 149 3. Additional surface detail

150 The heightfield simulation has a fixed size resolution which  
151 imposes a limit on the detail scale that can be achieved. We can  
152 add other simulations that improve the details by changing the  
153 surface normals locally. Furthermore, other advected properties  
154 as, e.g., surface foam, can also be simulated and applied to the  
155 final visualization. This section will introduce how these effects  
156 are accomplished.

### 157 3.1. Lower scale detail

158 From the heightfield of the fluid surface we can extract ap-  
159 propriate normals although they are restricted to the simulation  
160 resolution. We can increase the detail of the fluid just using nor-  
161 mal mapping. For example, [2] applied a normal map texture  
162 generated from the FFT ocean simulation by [6] and advected  
163 it as in [9].

164 The FFT ocean simulation from [6] is based on the com-  
165 putation of the Fourier amplitudes of a wave field. The final  
166 heightfield is obtained from the inverse FFT to those ampli-  
167 tudes. In our case, we compute the FFT each frame and obtain  
168 a normal map from its heightfield which is then applied to the  
169 fluid surface, as can be seen in Figure 4a.

170 An alternative to the FFT approach is the use of noise tex-  
171 tures with the same goal at mind. We can use gradient noise,  
172 being Perlin noise [28] the more popular, to obtain heightfields  
173 and compute normal maps from them to apply to the fluid sur-  
174 face. With 3D noise, we can create the illusion of animation  
175 moving through one of the dimensions. However, noise tex-  
176 tures have some inherent problems: it is not clear how to create  
177 a good water-like function and if tiling is required, the pattern  
178 repetitions are quite obvious, as shown in Figure 4b.

### 179 3.2. Surface Foam

180 In the real life situation where splashes are generated, like  
181 in breaking waves, it is most probable that foam is generated  
182 when these splashes hit the fluid bulk again.

183 In contrast to [2], where diffuse disks are generated and ad-  
184 vected with the fluid when particles fall into the surface fluid  
185 again, we simplify the idea. Using a floating-point single com-  
186 ponent texture mapped to the surface fluid, we detect where a




Figure 3: Breaking wave example from [1].




Figure 5: Foam is generated at particle-surface hit points and advected in successive frames using the fluid’s velocity.

187 particle has fallen and initialize that texel to a certain maximum  
 188 time-to-live (TTL) for the foam. This texture is then advected  
 189 using the fluid’s velocity field, tracing back as in [29]. Each  
 190 frame, the values of the texture are decreased  $\Delta t$  until they be-  
 191 come 0. These values are then multiplied with the desired foam  
 192 color and mapped to the fluid mesh, resulting in the blended  
 193 foam.

194 Using a texture for the foam introduces a constraint, how-  
 195 ever: its resolution should be dictated by the size of the particles  
 196 as, it could happen that more than one texel should be initial-  
 197 ized, depending on the particle to texel size ratio or, conversely,  
 198 that the texels are too big for the particle size.

199 Overall, as seen in Figure 5, the results are convincing and  
 200 the computations are faster due to the limited requirements,  
 201 which make it ideal in real-time applications.

#### 202 4. Photon-based Caustics

203 In order to add caustics to our real-time fluid simulation,  
 204 we follow the same path of [3] and extend their work. They  
 205 raycast a grid of photons, as points, through the scene in an  
 206 orthographic space defined at the light source, which also al-  
 207 lows them to easily add shadow mapping. One restriction they  
 208 have is that the depth map used in the raycast phase should be  
 209 continuous or, at least, with no great jumps. Other limitations  
 210 this technique has are the same as image-based rendering: the  
 211 results depend on the resolution of the textures used, which in  
 212 this case restricts where the photons can end within the scene.

213 Our contributions to their algorithm imply extending the  
 214 raycast of photons out of the light space to screen space, splat-

215 ting them oriented with the surface of the receiving mesh. In  
 216 contrast to [3], we do not generate a caustics map; the splats are  
 217 blended with the scene, varying their intensity depending on  
 218 the orientations of the caustic generating fluid position, as well  
 219 as the distance the photon has travelled inside the fluid. As our  
 220 fluid simulation is represented just by its surface, we restrict our  
 221 approach to refracted photons which will fall in the underlying  
 222 terrain and ignore reflected ones.

223 The multipass algorithm can be summarized in the follow-  
 224 ing steps:

- 225 1. Render the objects of the scene (excluding the fluid) from  
 226 the camera and store depth and normal maps.
- 227 2. Render the objects of the scene (excluding the fluid) from  
 228 light with an orthographic projection and store the depth  
 229 map.
- 230 3. Render the fluid from light with the same orthographic  
 231 projection as before and store the world positions and re-  
 232 fracted directions at each pixel.
- 233 4. Render the grid of photons. The primitives used are points  
 234 which will be expanded to quadrilaterals when a final po-  
 235 sition is found.

236 This grid of points has the same resolution as the ortho-  
 237 graphic projection used previously in Step 2. In a vertex shader,  
 238 the vertices will be raycast first in light space using the depth  
 239 map from Step 2. If there is no intersection found, i.e., the pho-  
 240 ton exited through a wall of the frustum, the raycasting will be  
 241 repeated in camera space. If there is not an intersection yet, the  
 242 point is discarded (rendered out of frustum). Otherwise, if an  
 243 intersection is found at light space, it is transformed to camera  
 244 space and checked for correctness:

- 245 • If the point is occluded in camera space, it is discarded.
- 246 • Else, if the point is not occluded and the depth does not  
 247 match between light and camera spaces, the raycasting  
 248 continues from the actual point position in camera space.
- 249 • Else, the point is correct, that is, the depths match be-  
 250 tween light and camera spaces, thus the point is final.

251 When a final point is found, from the previous condition  
 252 or from the camera raycasting, the normal is looked up in the  
 253 normal map from Step 1. In a geometry shader, the points are  
 254 expanded to quads oriented with their associated normal. Fi-  
 255 nally, in the fragment shader, the photons are textured with a

256 Gaussian splat, and their intensity is regulated depending on  
 257 how they are facing the light and the distance they have trav-  
 258 elled through the fluid until finally hit the receiving surface. At  
 259 last, they are blended to the contents of the framebuffer.

260 We have not implemented shadows to keep the algorithm  
 261 simple but, as suggested in [3], shadow mapping is easily added  
 262 as the depth map from light is already stored for the raycasting.

263 As can be seen in Figure 6, the visual results are good enough  
 264 for real-time rendering and the photons are not restricted to the  
 265 light space. For a full physically-based render, the precise radi-  
 266 ance of the photons should be computed. In order to make the  
 267 algorithm more approachable, we just regulate the photons con-  
 268 tributions with their orientation and user parameters as they are  
 269 just blended with the framebuffer, which allow the technique to  
 270 be faster in comparison, because we don't need the expensive  
 271 operations for gathering the photons.

## 272 5. Screen-space Refraction and reflection

273 Similarly to the caustics approach, we implement refraction  
 274 and reflection raycasting through depth maps, in the same way  
 275 of [23].

276 For simplicity, we want to be able to use the same raycasting  
 277 algorithm for both refractions and reflections, so we improve  
 278 upon previous works by rendering in separate buffers what is  
 279 above and below the fluid. This allows to, using the same code,  
 280 just look for ray-depth intersection in the appropriate buffer to  
 281 obtain the result and do a final composition with both refrac-  
 282 tions and reflection as needed.

283 In this case, rays are cast from camera and reflected or re-  
 284 fracted (or both) when they hit the fluid, as shown in Figure 7.

285 Here, we also use a multipass algorithm that can be ex-  
 286 plained in the following steps, always rendering from camera:

- 287 1. Render the fluid mesh and store the depth buffer.
- 288 2. Using two render targets (RT) named 'over' and 'below',  
 289 which will store color and depth, we render the objects  
 290 of the scene (dynamic objects and ground in this case)  
 291 and compare the depth with the previously stored. If the  
 292 depth is greater, the fragment is stored in the 'below' RT,  
 293 otherwise in the 'over' one. This pass can be thought as  
 294 a stencil test, which separates what is above or below the  
 295 fluid surface.
- 296 3. Render the fluid again, using the RTs. For each fragment  
 297 of the fluid two rays are cast: one for refraction (using  
 298 RT 'below'), one for reflection (using RT 'over'). As  
 299 the rays start from the camera, if the reflected/refracted  
 300 rays should come back, they are discarded. The results  
 301 of both raycasts are combined using Schlick's approxi-  
 302 mation [30] to Fresnel terms for simplicity.
- 303 4. Finally, to avoid the repeated render of the other objects  
 304 of the scene, we just use a screen-sized quadrilateral tex-  
 305 tured with the color buffer from the 'over' RT.

To reduce somewhat the need of the double raycasting, we  
 can compute the Fresnel term from [30] prior to the raycastings  
 at Step 3 as

$$F = F_0 + (1 - F_0)(1 - \theta)^5, \quad (5)$$




Figure 6: Caustics in the ground below the fluid surface with planar (boat) and noisy (buoy) terrain.

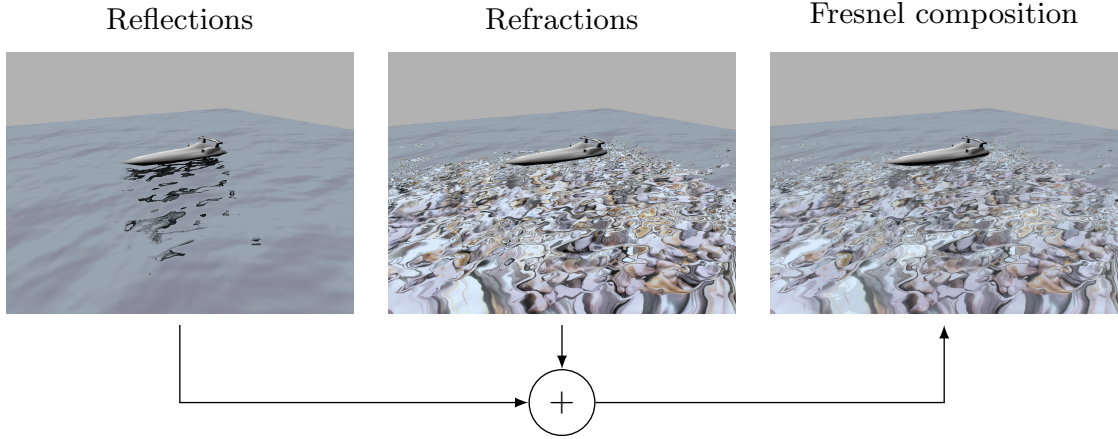


Figure 7: Reflections and refractions are found from the raycasting two different depth maps and finally composed using Fresnel for the final rendering.

306 being  $\theta$  half the angle between the ingoing and outgoing light  
 307 directions and  $F_0$  the known value of  $F$  when  $\theta = 0$ , the re-  
 308 flectance at normal incidence. As we use the value  $F$  for a  
 309 linear interpolation between the refracted and reflected colors,  
 310 we can impose a threshold  $\epsilon$  such as:

- 311 • If  $F < \epsilon$ , only the refraction raycasting is executed.
- 312 • If  $1 - F < \epsilon$ , only the reflection raycasting is done.
- 313 • Otherwise, both raycastings are done.

314 Additionally, for zones where the fluid height is quite low,  
 315 controlled by an user parameter, we get the color from the di-  
 316 rect view ray and interpolate from it to the combined color ob-  
 317 tained from the previous algorithm, using the depth difference  
 318 between the fluid and the ground below it. This alleviates some  
 319 visible artifacts caused by the triangular mesh used for the fluid  
 320 rendering, as shown in Figure 8.

321 Everything is done in screen-space, so there may be zones  
 322 where there is not enough information, i.e., a ray should hit a  
 323 point in space not visible; in those cases we detect the jump in  
 324 the depth map and make use of the last pixel with information  
 325 in the texture. Although this is really a problem due to lack  
 326 of information, it may remain greatly unnoticed with the ani-  
 327 mated lower scale detail techniques of Section 3.1 and the own  
 328 movement of the fluid surface.

## 329 6. Results and Discussion

330 We have tested the previous algorithms on an Intel Core2Duo  
 331 E8400 with 4GB of RAM and a Nvidia GTX280 running Ubuntu  
 332 11.10. The resulting averaged timings of the caustics and re-  
 333 fraction/reflection algorithms are shown in Table 1, as these are  
 334 the ones that tax more on the GPU by the use of raycasting.

335 An improvement to the normal mapping for lower scale de-  
 336 tail technique could be provided by also applying the technique  
 337 from [9], in which multiple sets of texture coordinates are used  
 338 and advected, already exploited in [2].




Figure 8: Artifacts from the fluid's triangular mesh on the left, alleviated on the right by interpolating the color value between the fluid's color and the ground color depending on the view distance from surface to ground.

Viewport Size	Caustics Resolution	Caustics	Refraction & Reflection
512 <sup>2</sup>	128 <sup>2</sup>	1.0058	2.74639
	256 <sup>2</sup>	2.35144	2.79409
	512 <sup>2</sup>	7.92134	3.01034
	1024 <sup>2</sup>	51.3408	3.17178
1024 <sup>2</sup>	128 <sup>2</sup>	1.2585	5.79334
	256 <sup>2</sup>	3.33537	5.98238
	512 <sup>2</sup>	12.8643	6.21948
	1024 <sup>2</sup>	70.9195	6.53224

Table 1: Averaged timings in milliseconds for frame for the caustics and refraction/reflection algorithms. The Viewport column indicates the viewport resolution. Similarly, the Caustics Resolution column indicates the size of the viewport used for the orthographic camera, and thus, the number of photons traced.

The foam simulation from [2] could solve the fixed-size texture restrictions of our current solution, as they simulate foam directly with advected diffuse disks on the fluid surface, although this comes at the additional cost of generating and maintaining these disks on the fly. An alternative we believe would help our foam simulation is the use of a pyramidal texture approach; when particles fall to the fluid they initialize the correct level of the pyramid, being the other levels initialized extrapolating from that one.

Nevertheless, the timing results for both these techniques combined, the lower scale detail and the foam advection, never exceed the 2ms mark.

For the caustics, as shown in Table 1 and concluded in [3], the performance of the caustics algorithm depends primarily on the size of the grid of photons, but in our case also on the direction of the light, which can cause more photons to miss the light space raycast and use the second camera space one, thus increasing the number of computations and texture fetches needed to try to find a final position for them. Also, as the raycasting is done in the vertex shader it is further slowed down because of the increased penalty of texture fetches in that shader stage. Although a direct comparison with [3] is difficult because of the different hardware used, they reported to achieve about 200fps with a 128<sup>2</sup> photon grid, which is the same that saying that each frame costs 5ms to compute. With newer hardware but the dual light and camera space raycasts we propose, the cost of computing caustics is, in our case, below 2ms for the same configuration.

As the number of photons is limited, there may be zones over or undersampled; a hierarchical solution could help to solve this problem as shown in [18, 31], but we would require that it remains highly dynamic, as we are addressing the visualization of a moving fluid. Another thing worth researching would be to extend these caustics, if possible, to volumetric ones as those in [32].

The performance of the refraction/reflection algorithm is quite variable, it depends on the size of the viewport as well as the coverage of the fluid in screen: the more visible pixels, the more rays are cast. For fair comparison, the results in Table 1

were captured with the fluid covering the whole viewport, and even in this case, the whole algorithm does not cost more than 10ms for a reasonably sized viewport. In perspective, [23] made total internal refraction available although without surface reflection which, in the best case, reported 138fps, i.e., 7.24ms per frame on a Nvidia 8800 GTX, using only one bounce for internal refraction on a viewport of 512<sup>2</sup>. Although our GPU is newer than theirs, in a similar scenario, we achieve less than half their time with both refraction and reflection.

Evidently, the restriction of the refraction/reflection algorithm being a screen-space technique limits how much information is available for such refractions and reflections. The simplest solution to this would be the use of environment maps, but, as the height of the fluid can be quite different across the domain, the position where the environment maps were generated would constraint, and even clip, possible geometry for correct refractions or reflections. Other alternatives should be considered to solve this limitation.

Both algorithms, caustics and refraction, may also suffer other performance penalties depending on the tessellation of the objects of the scene in question. This is due to the multipass character of the algorithms and the requirement of the rendering of the scene to obtain the depth maps for later raycasting. Although we have not encountered this problem in our tests due to low polygonal complexity, it should be worth having in mind. As a note, the boat model has 300 triangles, the buoy has 11k, the dolphin has 4k and the fluid and the ground have 32k triangles each.

Finally, the particles have just been rendered as billboards using depth and normal replacement with a sphere model. As they represent splashes, we want to maintain their crisp representation so, to improve their appeal, some additional tweaking could be done as applying some noise to their normals or deforming them in the direction they are moving to simulate some motion blur.

Finally, we have only taken into account the visualization of the surface of the fluid, as shown in Figure 9, in the future it should be also a key point to research water rendering as in, e.g., [33], in order to provide a full featured visualization.

## 7. Conclusions

In this paper we have presented a full pipeline of different algorithms for the rendering of heightfield-based fluid simulations coupled with particles, although the different parts can be applied to other situations as well.

The complex light-related effects like caustics, refractions and reflections have been addressed using raycasting techniques which ensure a more realistic simulation and the constraint of the algorithms to be in screen-space keeps the quantity of memory used low enough.

Additionally we have applied foam and lower-scale detail by applying textures to the fluid mesh; techniques which are very low demanding in comparison to the previous ones and really help to enhance the final result.




Figure 9: A dolphin underwater. Caustics are generated and projected on the dolphin and the terrain, visible from the surface.

## 431 Acknowledgements

432 We would like to thank Pere-Pau Vázquez for his kind com-  
 433 ments on the preparation of this work. With the support of the  
 434 Research Project TIN2010-20590-C02-01 of the Spanish Gov-  
 435 ernment.

- 436 [1] Ojeda J, Susín A. Hybrid Particle Lattice Boltzmann Shallow Water for  
 437 interactive fluid simulations. In: 8th International Conference on Com-  
 438 puter Graphics Theory and Applications. GRAPP'13; 2013, p. 217–26.
- 439 [2] Chentanez N, Müller M. Real-time simulation of large bodies of water  
 440 with small scale details. In: Proc. ACM SIGGRAPH/Eurographics Sym-  
 441 posium on Computer Animation (SCA). 2010, p. 197–206.
- 442 [3] Shah MA, Kontinen J, Pattanaik S. Caustics mapping: An image-space  
 443 technique for real-time caustics. IEEE Transactions on Visualization and  
 444 Computer Graphics 2007;13(2):272–80.
- 445 [4] Bridson R. Fluid Simulation for Computer Graphics. AK Peters; 2008.
- 446 [5] Solenthaler B, Pajarola R. Predictive-corrective incompressible sph.  
 447 ACM Trans Graph 2009;28(3).
- 448 [6] Tessendorf J. Simulating ocean water. In: SIGGRAPH 2001 Course  
 449 Notes. 2001..
- 450 [7] Yuksel C, House DH, Keyser J. Wave particles. ACM Trans Graph  
 451 2007;26(3).
- 452 [8] Layton AT, van de Panne M. A numerically efficient and stable algorithm  
 453 for animating water waves. The Visual Computer 2002;18(1):41–53.
- 454 [9] Max N, Becker B. Flow Visualization Using Moving Textures. In: Pro-  
 455 ceedings of the ICAS/LaRC Symposium on Visualizing Time-Varying  
 456 Data. 1996, p. 77–87.
- 457 [10] Kajiyá JT. The rendering equation. In: Proceedings of the 13th an-  
 458 nual conference on Computer graphics and interactive techniques. SIG-  
 459 GRAPH '86; 1986, p. 143–50.
- 460 [11] Veach E, Guibas LJ. Metropolis light transport. In: Proceedings of the  
 461 24th annual conference on Computer graphics and interactive techniques.  
 462 SIGGRAPH '97; 1997, p. 65–76.
- 463 [12] Jensen HW, Christensen PH, Kato T, Suykens F. A practical guide to  
 464 global illumination using photon mapping. In: SIGGRAPH 2002 Course  
 465 Notes. 2002..
- 466 [13] Purcell TJ, Buck I, Mark WR, Hanrahan P. Ray tracing on programmable  
 467 graphics hardware. In: Proceedings of the 29th annual conference on  
 468 Computer graphics and interactive techniques. SIGGRAPH '02; 2002, p.  
 469 703–12.
- 470 [14] Nvidia . NVIDIA OptiX Application Acceleration Engine. <http://www.nvidia.com/object/optix.html>; 2012. [Online  
 471 as of June-2013].
- 472 [15] Stam J. Random caustics: natural textures and wave theory revisited. In:  
 473 ACM SIGGRAPH 96 Visual Proceedings: The art and interdisciplinary  
 474 programs of SIGGRAPH '96. SIGGRAPH '96; 1996..
- 475 [16] Szirmay-Kalos L, Aszódi B, Lazányi I, Premecz M. Approximate ray-  
 476 tracing on the gpu with distance impostors. Computer Graphics Forum  
 477 2005;24(3):695–704.
- 478 [17] Wyman C, Davis S. Interactive image-space techniques for approximat-  
 479 ing caustics. In: Proceedings of the 2006 symposium on Interactive 3D  
 480 graphics and games. I3D '06; 2006, p. 153–60.
- 481 [18] Wyman C. Hierarchical caustic maps. In: Proceedings of the 2008 sym-  
 482 posium on Interactive 3D graphics and games. I3D '08; 2008, p. 163–71.
- 483 [19] Guardado J, Sánchez-Crespo D. Rendering water caustics. In: GPU  
 484 Gems. Addison-Wesley; 2004, p. 31–44.
- 485 [20] Yuksel C, Keyser J. Fast Real-time Caustics from Height Fields. The  
 486 Visual Computer (Proceedings of CGI 2009) 2009;25(5-7):559–64.
- 487 [21] Sousa T. Generic refraction simulation. In: GPU Gems 2. Addison-  
 488 Wesley; 2005, p. 295–305.
- 489 [22] Wyman C. An approximate image-space approach for interactive refrac-  
 490 tion. In: ACM SIGGRAPH 2005 Papers. SIGGRAPH '05; 2005, p.  
 491 1050–3.
- 492 [23] Davis ST, Wyman C. Interactive refractions with total internal reflection.  
 493 In: Proceedings of Graphics Interface 2007. GI '07; 2007, p. 185–90.
- 494 [24] Hu W, Qin K. Interactive approximate rendering of reflections, refractions,  
 495 and caustics. IEEE Transactions on Visualization and Computer  
 496 Graphics 2007;13(1):46–57.
- 497 [25] Salmon R. The lattice boltzmann method as a basis for ocean circulation  
 498 modeling. Journal of Marine Research 1999;57(3):503–35.
- 499



- 500 [26] Schaufler G. Nailboards: A rendering primitive for image caching in  
501 dynamic scenes. In: Proceedings of the Eurographics Workshop on Ren-  
502 dering Techniques '97. 1997, p. 151–62.
- 503 [27] van der Laan WJ, Green S, Sainz M. Screen space fluid rendering with  
504 curvature flow. In: Proceedings of the 2009 symposium on Interactive 3D  
505 graphics and games. I3D '09; 2009, p. 91–8.
- 506 [28] Perlin K. Improving noise. In: Proceedings of the 29th annual conference  
507 on Computer graphics and interactive techniques. SIGGRAPH '02; 2002,  
508 p. 681–2.
- 509 [29] Stam J. Stable fluids. In: Proceedings of the 26th annual conference on  
510 Computer graphics and interactive techniques. SIGGRAPH '99; 1999, p.  
511 121–8.
- 512 [30] Schlick C. An Inexpensive BRDF Model for Physically-based Rendering.  
513 Computer Graphics Forum 1994;13(3):233–46.
- 514 [31] Wyman C, Nichols G. Adaptive caustic maps using deferred shading.  
515 Computer Graphics Forum 2009;28(2):309–18.
- 516 [32] Liktov G, Dachsbacher C. Real-time volume caustics with adaptive beam  
517 tracing. In: Symposium on Interactive 3D Graphics and Games. I3D '11;  
518 2011, p. 47–54.
- 519 [33] Gutierrez D, Seron FJ, Munoz A, Anson O. Visualizing underwater ocean  
520 optics. Computer Graphics Forum 2008;27(2):547–56.