Aalborg University Copenhagen

Department of Medialogy

Semester: Medialogy 10

Title: Optimizing water simulation for visual effects through perception studies.

Project Period: Feb 1st 2011 – Oct 3rd 2011

Semester Theme: Master Thesis

Supervisor: Rama Hoetzlein

Aalborg University Copenhagen
Lautrupvang 15, 2750 Ballerup, Denmark

Semestercoordinator: Stefania Serafin
Secretary: Lisbeth Nykjær
Phone: 9940 2471
lny@create.aau.dk
https://internal.media.aau.dk/

Created by:

Anders Heilemann

Copies: 3

Pages: 46

Finished: October 3rd

Abstract

This thesis investigates if the presence of objects related to the story of a video affects the perceived quality of a background environment, such as a water simulation. The aim was to provide guidelines for small- to mid-sized visual effects studios, in order for them to optimize the performance of their water simulations.

A theory was created based on the existing theories of inattentional blindness and guided search, stating that the more important an object in the scene was to the story, the more attention it would receive. Consequently the environment around the elements important to the story would receive less attention.

Four videos with the same scene setup with a river flowing towards a waterfall were created; two in low simulation quality and two in high. One of each had a boat flowing down the river. The theory stated that in the videos with the boat, the test subjects would not notice the quality of the water as much.

The results partially confirmed the theory, the test subjects rated the quality of the water equal in the low quality version with and without the boat. However the majority of test subjects commented that the presence of the boat had an influence on their ability to remember the quality of the boat.

The thesis concludes that inattentional blindness is indeed relevant for the production of computer graphics, however to know exactly its impact, more investigations still need to be made.

Keywords: Hydrodynamics, water simulation, perception, inattentional blindness, guided search, Naiad, Houdini.

PREFACE

This report was written at Aalborg University Copenhagen as a result of the work done by Anders Heilemann while attending the 10th semester of Medialogy.

The referencing system for citing works in the report follows the guidelines from APA 6th Edition.

Following this report is a disc located at the back containing the digital equivalent of the report, together with any media created during the implementation phase of the report.

The author would like to thank Exotic Matter for their invaluable sample projects, as well as the members of their support forum. They served as an invaluable asset in the implementation phase of this project.

CONTENTS

LIST OF FIGURES

INTRODUCTION

In the context of creating visual effects for films, the rise in demand for even higher degrees of realism increase each time a new computer graphics film is released. Mike Mcgree, one of the founders of the visual effects company Framestore, once said in an interview with fxguide.com

> PEOPLE ARE ALWAYS SAYING, 'CAN YOU DO SOMETHING WE'VE NEVER SEEN BEFORE?' TO VISUAL EFFECTS ARTISTS. (I. FAILES, 2011B)

Volker Engel and Marc Weigert, the co-producers and visual effects supervisors from Digital Domain, talks about the 'Decrease in Cost' vs. 'Increase in Expectation', and how new hardware and software makes creating effects cheaper while experience makes it easier. But as the cost goes down the expectations go up, so the studio always needs to reinvent and create effects that no one have ever seen before, which stretches the software and the computing power to its absolute limits. They end with a statistic that says that rendering the LA Earthquake scene of the motion picture 2012 would take 16 years on a single machine (Verrier, 2011).

This is just some of many examples supporting that the visual effects studios need to outshine each other to be the best on the market when it comes to creating stunning and realistic scenes for the increasingly demanding audience. However as this demand rise, the studios' need for computational power, rise as well.

Several small as well as big studios around the world close down as they have to take on challenges they might not be fit to complete on budget. At the same time they have to underbid the rising competition from eastern countries where salaries are but a fraction of the rest of the world. In California which for many years has been the capital of visual effects creation, most companies operate with profit margins less than 5% (Verrier, 2011).

The two most time and money demanding areas of creating computer graphics are the rendering of images and the simulation of fluids. Rendering is one of the final stages of creating an image where light, shadows and other effects are calculated. Fluid simulations cover the creation and animation of everything which cannot be defined as solid, such as water, smoke, fire and similar effects. This thesis will focus solely on the simulation of water.

The motion picture "2012" from 2009 directed by Roland Emmerich tells a story about the sun heating up the earth from the inside, resulting in the earth crust breaking apart causing massive tsunamis and havoc throughout the world. The film showed some of the most complex water simulations created. In an interview with the computer hardware manufacturer "Nvidia" Stephan Trojansky, the president of the visual effects company Scanline, said that because of the sheer size of their simulations and the detail level, they created 1.2 petabytes of data during their 103 shots for the film (Priscaro, 2009). The cost to just gain access to the hardware needed to work with simulations of such size and complexity would make it difficult for many studios to compete. But is the increased effort and complexity really worth it?

"Pirates of the Caribbean: At World's End" released in 2007 and directed by Gore Verbinski is the third film in a series of motion pictures where the general story is about the journeys of the pirate "Captain Jack Sparrow". It featured a very dramatic scene where two ships fire upon each other while sailing around and down a maelstrom, meanwhile the captains of the two ships have a sword fight on the spars of one of the ships. The water simulation in the background is impressive, but only really visible when the video is set on pause because of the heavy motion blur caused by fast camera movement. Examples of the maelstrom scene can be seen in Figure 1; the right image especially shows how an otherwise impressive water simulation is sent to the background and becomes quite difficult to see, due to the actions of the main characters in the foreground.

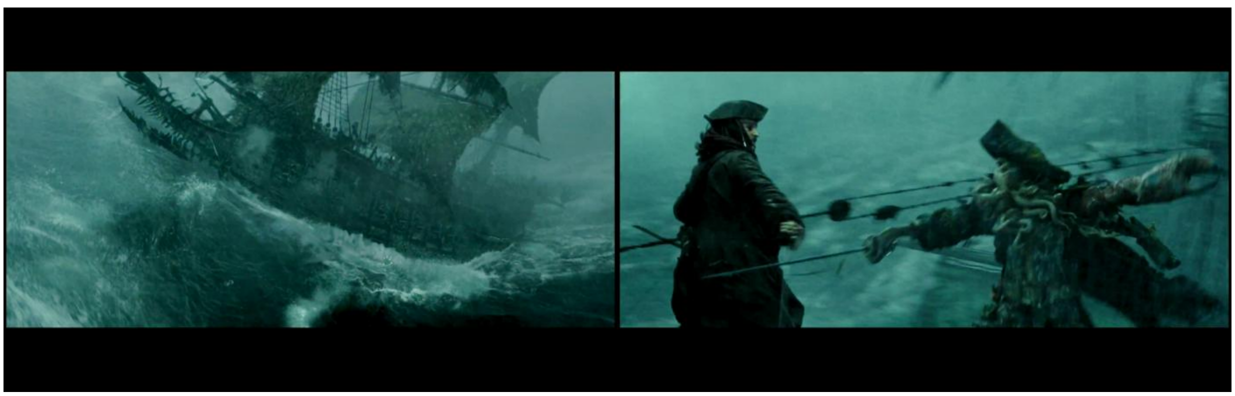

FIGURE 1 - SCREENSHOTS FROM "PIRATES OF THE CARIBBEAN: AT WORLDS END" – WITH HEIGHTENED BRIGHTNESS TO INCREASE VISIBILITY ON PRINT. (VERBINSKI,G. 2007)

While the demand for complexity and realism increases, research on whether the audience actually perceives the high quality when it is set in context, such as in a motion picture, is scarce.

Bordwell suggests that when an audience watches a motion picture, they tend to focus on the objects related to the story, he quotes:

> SINCE STORY COMPREHENSION IS ONE OF OUR PRIMARY TASKS IN WATCHING A MAINSTREAM MOVIE, WE WILL TEND TO IGNORE OTHER THINGS. (BORDWELL, 2011)

Consequently in the example from above, the audience will tend to look at the sword fight and not at the water simulation, as the sword fight is more important to the story.

This thesis will attempt to create a video with a water simulation of similar detail level to those used in motion pictures, but in a smaller form and in a more controlled environment. Through a user study it is then the aim to outline which parameters of the simulation are the most important for increasing the perceived quality of water simulations, and what their effect is on the computational requirement. The purpose is to investigate to what degree the typical audience notice quality changes as the detail level of the water simulation changes, when it is implemented in a story context.

The target of this research is small- to mid-sized visual effects studios, in order to provide them with guidelines as to how low they can set their simulation quality, while still delivering satisfying results to their target audience. By lowering the simulation quality, the visual effects studios can cut down on the computation requirement, and thereby save time and money, increasing their competitive edge against other studios.

The rest of the thesis will take basis in the thesis statement below:

To what degree is the perceived quality of a water simulation affected by the presence of objects related to the story in a video.

To answer this thesis statement, an analysis chapter will investigate two areas. First how the human visual system interprets images and why viewers tend to ignore objects not related to the story in a video. The second part investigates how to create water simulations to aid the implementation process of the thesis. Based on the analysis a test method will be formulated, and based on this method a water simulation will be designed, implemented and tested against the thesis statement. Finally the results of the testing will be evaluated and the thesis statement answered followed by a retrospective discussion on the results and the process.

ANALYSIS

This chapter is divided into two parts; the first part investigates visual perception theories and relates them to the quality perception of water simulations. The second part investigates the basics of water simulations to gain an understanding of how to create simulations of a similar quality as those implemented in motion pictures. Merging this knowledge can then be used to design a solution testable against the thesis statement.

VISUAL PERCEPTION

The idea to design computer graphics based on human perception is not new. During the last 50 years an extensive amount of research has been made in visual perception, where the focus has been trying to understand how the eyes and the brain perceive and interpret images. The last decade much of this research has been combined with computer graphics.

This chapter investigates the visual perception theories of inattentional blindness and change blindness, two theories attempting to explain why viewers do not remember all the details of an image. Following is an in-depth investigation presenting what pre-attentive processes make those theories occur, as well as how others have used this knowledge for optimizing rendering in computer graphics. The goal is to gain knowledge of how audiences perceive images in film and how to design accordingly.

CHANGE BLINDNESS

Change Blindness is the study of how viewers are unable to detect changes in a scene, even after having been exposed for several seconds. It mostly occurs when the viewer has not abstracted certain details of a scene through direct attention. Simons made an overview of five different change blindness theories, shown in Figure 2, the images show what is remembered after an image sequence has been displayed (D. Simons, 2000).

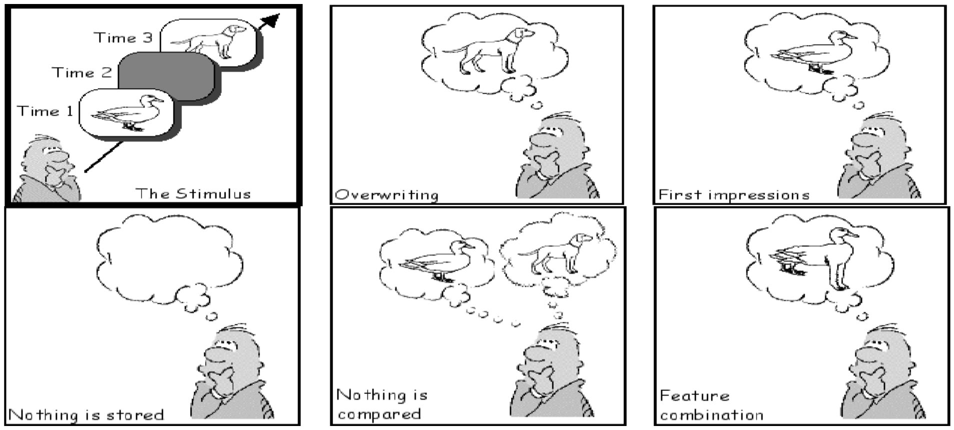

FIGURE 2 – IMAGE SERIES FROM SIMONS VISUALIZING THE RECENT CHANGE BLINDNESS THEORIES (SIMONS,DANIEL 2000)

To transfer the theories to the perceived quality a viewer has of a reappearing computer graphics object, such as a water simulation, the stimulus order would be every time the same water simulation reappears. In the first theory, **overwriting**, the viewer forgets the details of the water as soon as it reappears; consequently it is only important to display the water in high quality at the last appearance. Number two is **first impressions** where the quality of the water will only be processed the first time, subsequent appearances are then retrieved from memory. Next up is **nothing is stored** where the details of the water are only remembered if it has received direct attention. Number four is **nothing is compared** and states that each appearance is remembered individually. The last theory **feature combination** suggests that all appearances are combined to a single abstraction of the quality (Healey & Enns, 2011).

Set in relation to this project, then change blindness requires the water to reappear in successive scenes. It could be tested if detail degradation between cut-scenes would be noticed; i.e. if the water is first be presented in high quality and subsequently in low quality will the audience then notice the quality drop. However this would be more relevant for rendering where each frame or sequence can be rendered with independent settings. In simulations each frame is dependent on the frame before it, requiring different simulations for each cut scene; which would be tedious and make it difficult to adapt to changes from to a director.

The change blindness theory **nothing is stored** is however different than the others, it does not require reappearances in successive scenes; it simply states that nothing will be remembered unless it receives dedicated attention. Very similar to that change blindness theory is the theory of inattentional blindness which investigates the same phenomenon that Bordwell also talked about, namely that humans tend to ignore things that are not part of the primary task.

INATTENTIONAL BLINDNESS

Inattentional blindness is the theory that otherwise salient stimuli will go unnoticed, even when located in the fixation area. The reason for such conspicuous elements to be ignored, although they would normally attract attention is that they do not support the task at hand. The first studies suggesting this was made by Yarbus who pioneered the study of saccadic eye movement by using eye tracking, here test subjects were given different tasks before looking at an image. Figure 3 shows an example of this where one test subject was told to guess the ages of the people in the image (left) and another was told to remember the position of objects and characters in the image (right). It clearly shows that depending on the task at hand the test subjects had very different viewing patterns (Yarbus, 1967).

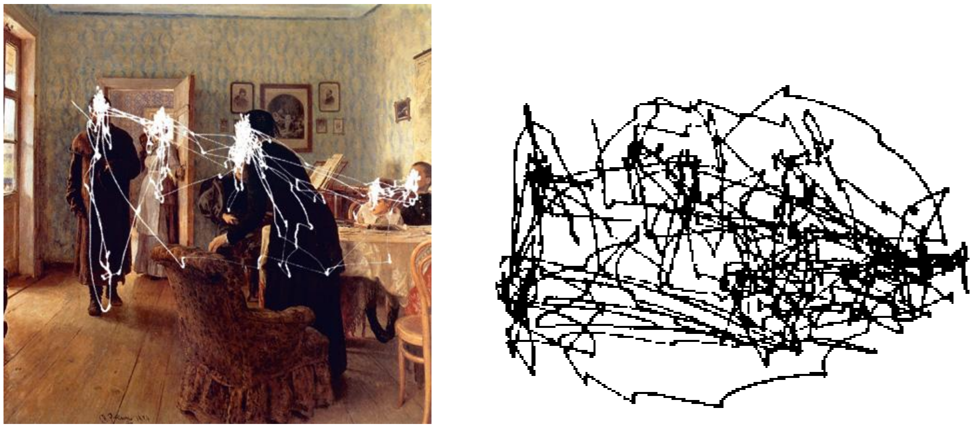

FIGURE 3 – IMAGES FROM YARBUS EYE TRACKING EXPERIMENTS, SHOWING AN EXAMPLE OF THE EYE MOVEMENT FROM A TEST SUBJECT WITH THE TASK TO GUESS THE CHARACTERS AGE (LEFT) AND WITH THE TASK TO REMEMBER POSITIONS OF CHARACTERS AND OBJECTS (RIGHT) (YARBUS, 1967).

For inattentional blindness to occur it is a condition that the observer follows a task (K. Cater, Chalmers, & Ledda, 2002). Simons et al. did an experiment where viewers completed the task of counting basketball passes the white team made, because they were so focused on the task most test subjects did not notice a man dressed as a bear walking through the scene, as seen in the image from their testing in Figure 4 (D. J. Simons & Chabris, 1999). This experiment attracted attention and has consequently been repeated numerous times, for example for a series of commercials made by Transport for London where they use the phrase “It’s easy to miss something you’re not looking for” to tell car drivers to pay attention to cyclists (Transport for London, 2009).

FIGURE 4 – IMAGE FROM SIMONS INATTENTIONAL BLINDNESS EXPERIMENT, SHOWING A MAN IN A GORILLA SUIT (SIMONS,D J. 1999)

These latter examples suggests that if an audience can be this blind to objects in the scene because they follow a task, and that the task can be story comprehension as Bordwell suggests. Then the quality of a water simulation as the one from Pirates of the Carribean exemplified in Figure 1 will not be abstracted or remembered.

Inattentional blindness therefore suggests that the simulation detail might as well be low from the beginning. In order to comprehend when the quality can be lowered and when it cannot, it is important to understand how exactly the audience perceives images; and what they dedicate attention to. This next chapter presents an in-depth view into some of the underlying preattentive processes that cause inattentional blindness to occur. The purpose is to understand how to design a scene according to the audience's perception

PREATTENTIVE PROCESSING

In order to design a product capable of being tested against the effect of inattentional blindness, this chapter investigates the preattentive processes that cause inattentional blindness to occur. Healey expresses why this is important with this quote:

> MANY VIEWERS ASSUME THAT AS WE LOOK AROUND US WE ARE CONSTRUCTING A HIGH-RESOLUTION, FULLY DETAILED DESCRIPTION OF WHAT WE SEE. RESEARCHERS IN PSYCHOPHYSICS HAVE KNOWN FOR SOME TIME THAT THIS IS NOT TRUE (HEALEY & ENNS, 2011).

To learn what triggers inattentional blindness it is necessary to understand why that is not true and what then actually happens. This knowledge of how the human visual system processes what it sees can be used to create guidelines for designing a product with high perceived quality using low computational effort.

The human visual system is only capable of decoding shape and color from a very small region at any time. In order to see detailed information from more regions, the eyes flicker rapidly between regions where they focus for a short duration. The focus period is called a fixation and shift of focus is called a saccade, this fixation-saccade cycle happens 3-4 times every second. The task and goals of our current mental state guides the saccades in a top-down fashion. This is essentially the cause of inattentional blindness; the eyes just cannot attend all objects simultaneously. What the visual system chooses to fixate upon is determined almost instantaneously by the preattentive vision; consequently in order to design for inattentional blindness it is necessary to understand how the preattentive vision prioritizes what to guide the eyes towards.

Exactly how the preattentive vision works is still a subject of discussion, one of the first theories was the feature integration theory by Treisman which stated that specific features can be identified preattentively, like a red object in a scene of blue objects, or a curved line between straight lines, but in order to combine these features a fixation was required. So if the task is to find a something with a unique feature, the task can be completed almost instantaneous as the preattentive searches occur in parallel. If the task was to locate a combination of different features a serial search through the objects of the scene would be needed, requiring multiple fixations. The more distinct the features of the target are, compared to the distractors, the less time it takes to complete the search (Healey & Enns, 2011; Treisman, 1985). The Feature integration theory is based on the dichotomy that serial and parallel processes were completely separate. However some researchers did not support this dichotomy but instead believed that groups of neurons in the brain were competing over time to represent the same objects. Quinlan, Humphreys and Duncan did several studies on what they called the Similarity Theory and hypothesized that search time was based on two criteria: T-N Similarity, meaning the amount of similarity between the target and nontargets and N-N Similarity meaning the amount of similarity between the nontargets. Essentially it means that the more distinct the target is compared to the rest of the image, the less time it takes to find it because a more preattentive and parallel process can be used (Healey & Enns, 2011).

Figure 5 illustrates the similarity theory where viewers tend to spent more time looking at all the objects in the image before finding the target L shape in image (b) compared to image (a).

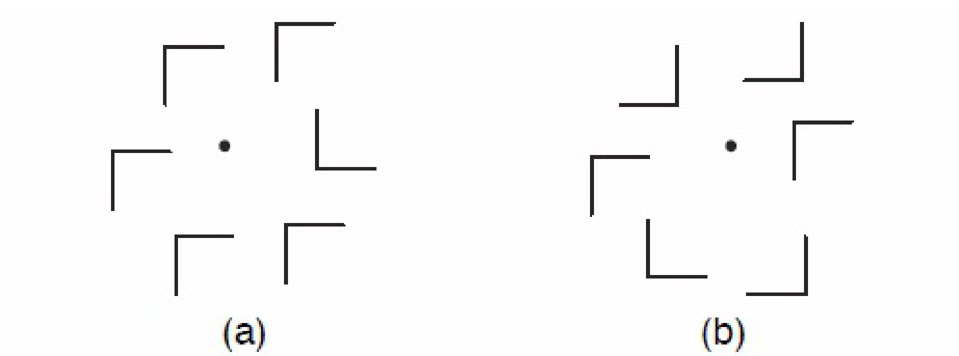

FIGURE 5 - EXAMPLE IMAGES FROM DUNCAN AND HUMPHREYS STUDIES ON THE SIMILARITY THEORY - BOTH (A) AND (B) CONTAINS ONLY 1 L SHAPED OBJECT, IT IS HOWEVER FASTER TO IDENTIFY IN (A) BECAUSE THE NON-TARGET SIMILARITY IS HIGHER. (J. & W., 1989)

Setting this in context with following the story of a motion picture, the faster the audience can find the targets most important for the story in the motion picture, the less time they use to attend other objects in the scene; such as the environment. This also explains why test subjects in Simons Inattentional Blindness experiment did not see the bear, their task was to count the passes the white team made, thereby preattentively sorting away all dark targets, such as the bear.

Wolfe et al. were the first to suggest that both bottom-up and top-down processes were used during visual search. With their Guided Search theory, they hypothesized that an activation map was created preattentively. The level of activation at each location on the map would reflect the likelihood of the target being at that location, i.e. how important the object is to the story. The level of activation would be distributed according to how many features were shared between the object and the target. Dedicated attention would then be distributed across locations according to the activation level. For example, if the target is a green bar with a steep orientation, the preattentive processing would activate all objects with the green and steep feature; attention would then be distributed to the objects most activated, Figure 6 illustrated this example. This also supports the Similarity Theory in that search times increase when the target and non-target share more features. (Healey & Enns, 2011)

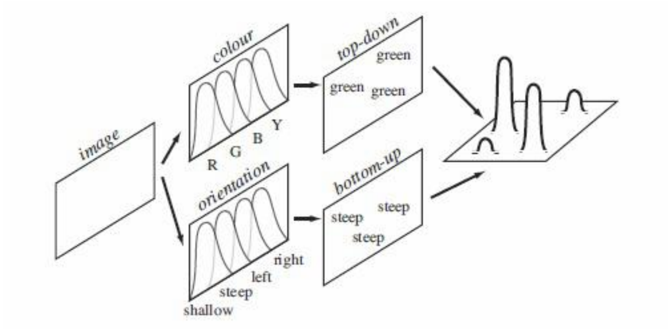

FIGURE 6 – IMAGE FROM WOLFE'S WORK ON THE GUIDED SEARCH THEORY SHOWING AN EXAMPLE OF AN ACTIVATION MAP (HEALEY & ENNS, 2011)

If the activation level of the audience could be estimated throughout a motion picture based on which object are the most important to the story, and how easy they are to locate whenever there is a change of camera, then production time could be distributed accordingly. Examples of optimizing rendering efforts based on visual perception show how others have had similar ideas.

RELATED WORK

Because realistic rendering of complex scenes is very computational expensive the theory of inattentional blindness has been used in selective rendering systems to increase the rendering speed, without decreasing the perceived quality of the rendering. As the theory states that when a task is given, the viewers would look mostly at the areas and objects related to that task, consequently the rendering effort in other areas can be degraded unnoticed.

Cater et al. found that depending on the task given, the saccade patterns of the test subjects changed accordingly, spending almost no time attending regions of the image unimportant to the task at hand. Consequently viewers did not notice when the rendering quality was lowered in those regions (K. Cater, Chalmers, & Ward, 2003).

Results from Sundstedt et al. showed similar results, but however found that this was the case both for their test group, but also for the control group who were given no task. Through eye-tracking they could see that all test subjects primarily attended the salient objects in the scene. Their results showed that they could decrease rendering times seven fold, without their test subjects noticing any particular drop in quality (Sundstedt, Debattista, Longhurst, Chalmers, & Troscianko, 2005).

Both Sundstedt et al. and Cater et al. used the knowledge that the human visual system uses two processes when looking at an image, to guide their rendering efforts. The first is a bottom-up process which is an automatic stimulus that registers elements in the image like; contrast, size, color, brightness, orientation, edges and motion. The second process is a top-down process which is voluntary and focuses on different elements depending on the viewer's task, i.e. Guided Search. Figure 7 from left to right show how Sundsted et al. created an importance map to guide their renderer. The task map from the top-down process (a+b), the saliency map from the bottom-up process (c), the maps combined to an "importance map" which guided the renderer (d) and lastly the final rendered image (e).

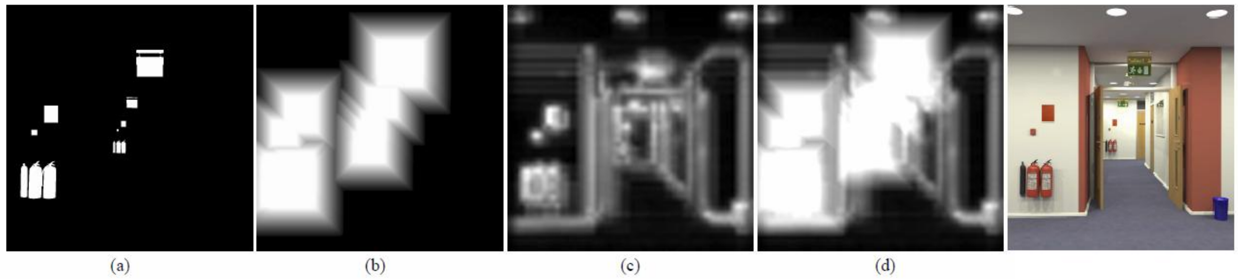

FIGURE 7 – IMAGES FROM SUNDSTEDS SELECTIVE RENDERING EXPERIMENTS, SHOWING THE CREATION OF AN IMPORTANCE MAP (SUNDSTEDT,V. 2005)

SUMMARY

In a motion picture context where story is the driving task, Guided Search would guide the eyes towards the objects with highest relation to the story. As it is the director who guides the audience through the story, he knows which elements are the most important to the story at any point in time. By marking these, the post production team could go through the scene layout and distribute production time, based on the anticipated importance each element would have on the story. Objects with low importance should have less time dedicated for modeling, simulating, rendering etc. Even without the director, this could still be estimated based on the general saliency object have; humans over animals over objects over environment. The director's input would be needed when objects are given additional importance. In the Motion Picture "Aladdin" directed by Clements and Musker from 1992 (Musker & Clements, 1992), the audience knows that a special lamp has great importance in to the story; consequently the audience might assign a higher activation level towards all lamps in the film, especially the one with the genie. Another example could be that the Pixar lamp attracts more attention than a normal lamp because it is given personality.

The more important an element is, the longer it takes for the visual system to process it. A location might even captivate the audience's attention so much it completely prevents them from saccading to other locations with lower activation levels. One problem with assigning production time based on an importance map is that it will always just be an approximation of the audiences' actual activation map. Furthermore, story is likely not the only thing guiding the audience's attention, different groups of people might prioritize differently, a sailor probably finds water more interesting than most people.

The last part of the analysis investigates the basics of fluid simulations; this is to understand the different approaches of solving water simulations and to aid in choosing the most optimal of those for use in this project.

COMPUTATIONAL FLUID DYNAMICS

To be able to create a water simulation with the size and detail level required to incite an audience, an investigation into how such simulations are created is presented. In traditional computational fluid dynamics accuracy is often the most important end goal. In computer graphics this is different because realistic fluid simulations have many interacting forces that have little or no visual consequence, and therefore no importance to the final image. This chapter investigates the basics of fluid dynamics and only delves deeper into the parts that are relevant for creating water simulations for computer graphics.

Fundamentally all water simulations for computer graphics are based upon variations of the basic fluid flow equation, Navier-Strokes. The result of the equation is a velocity field which describes the way the fluid wants to move at a given point in space and time. The equation is shortly presented here in order to identify the parameters commonly used for simulations, for more information the reader is referred to the book "Fluid Dynamics for Computer Graphics" by Robert Bridson (Bridson, 2008). The basic fluid flow equation looks like this:

$$\frac{D\vec{u}}{Dt} + \frac{1}{\rho}\nabla p = \vec{g} + v\nabla \times \nabla\vec{u}$$

$$\nabla \times \vec{u} = 0$$

The symbols in the first equation represent the different physical parameters that are usually required for fluid simulations; the second equation is the incompressibility condition which prevents the fluid from changing volume. Here follows an explanation of what the symbols mean:

- ρ = Density – The weight of the fluid e.g. oil weights more than water.
- p = Pressure – The amount of force per unit area the fluid exerts on anything around it
- $\vec{g}$ = Body forces – All the forces that affects the fluid body, such as gravity, wind etc.
- v = Viscosity – How much the fluid resists while it flows e.g. honey has a high viscosity, water has low.
- $\vec{u}$ = 3D velocity vector – Which direction if the fluid heading.
- $\frac{D\vec{u}}{Dt}$ = Material derivative – The change of velocity over time.

The incompressibility condition is one of the ways the equation has been optimized for computer graphics, water does compress in reality; this is what allows you to hear underwater. However the impact the compression has on the general movement of the water is so small it can be ignored.

There exist different implementations for solving the above equation, these can be divided into 3 categories. Either they are implemented using a Eularian method, a Lagrangian method or a combination of those two.

The Lagrangian method defines the simulation as a particle system where each particle contains data such as position and velocity, the particles can as such be thought of as molecules of the water. At every timestep during the simulation each particle samples its nearest neighbors in order to compute its own velocity. The first Lagrangian methods used a grid attached to the surface of a fluid during the entire computational process, because the grid was fixed to the surface as a mesh the methods were not very successful at efficiently simulating fluids with tendencies to distort and break apart, such as turbulent water often does. Smoothed Particle Hydrodynamics (SPH) was a mesh-free implementation of the Lagrangian method. This approach made it very adaptable to fluid deformations, and thus good for water splashes. One major disadvantage of SPH fluids is that inside bodies of water, large amounts of data has to be processed from often nearly motionless particles, which impacts the speed of such simulations (Liu & Liu, 2003; Next Limit Technologies, 2010; Bridson, 2008).

The Eularian method defines the fluid through a grid which is fixed in space, as the fluid then flows through the cells, the information in each grid cell defining the fluid at that location is consequently updated. Before SPH gained ground, the Eularian grid methods were dominant in most areas of computational fluid hydrodynamics. The disadvantage of the Eularian grid method is that the fluid cannot travel outside the predefined grid. The grid can be scaled up at the beginning of the simulation to cover the fluid during the entire simulation, however this consequently leads to increased computation requirements as all grid cells need processing at every time step (Liu & Liu, 2003; Bridson, 2008).

A range of different combinations of the two methods exist, one of such is the Particle-in-Cell (PIC) where all the fluid data except advection is stored on an Eularian grid, and updated at every timestep by taking a weighted average of the surrounding particles. However PIC's continued averaging at each frame results in numerical diffusion. Brackbill et al. solved this diffusion problem with their Fluid-Implicit-Particle (FLIP) method where all data is stored on the particles, but as new velocities are calculated the data is moved to a temporary Eularian grid for computation. Following the solve, the new velocities are distributed to the particles which change position accordingly. Because all data is stored in the particles and only moved to the grid during computation; the grid can be adaptive. This means it can grow or shrink along with the fluid movement, removing the bounding issue normally associated with using Eularian grids. (Brackbill, Kothe, & Ruppel, 1988; Bridson & Yongning, 2005)

Jeff Lait, a senior mathematician at Side Effects Software, the company behind the visual effects software Houdini, compares the speed of the FLIP method against the SPH and Eularian methods.

He explains that one speed advantage in FLIP comes from its efficiency at creating stable simulations with few or no substeps. He makes a rough estimate that volume fluids need 10x substeps, SPH needs 100x substeps where FLIP might not need any substeps at all. (Lait, 2010).

The FLIP method appears to be the most efficient algorithm for creating large water simulations for motion pictures, as it can efficiently handle both large bodies of water as well as define small splashes without constricting the fluid movement to predefined bounds. Several of the most popular commercial software implementations for simulating fluids dynamics for computer graphics have recently begun to implement FLIP methods, either to complement existing SPH and Eularian implementations or to replace them. This has happened after Robert Bridson in 2005 expanded the original FLIP method to be incompressible, making it more efficient for creating fluid simulations for motion pictures (Bridson & Yongning, 2005). Examples of the software that now use the FLIP algorithm include RealFlow by Next Limit (Next Limit Technologies, 2010), Houdini by Side Effects Software (SideFX Software, 2011) and Naiad by Exotic Matter (Exotic Matter, 2011b; I. Failes, 2011a; I. Failes, 2011b). All of these have been used for several high-end water simulations in motion pictures. Naiad is the software going to be used for simulations in this project because of its efficiency and accessibility. The free version has no other limitations than a 30 day trial period, where RealFlow limits particle count to 100.000 in the free version. Compared to the current version of Houdini; Houdini 11, the solver in Naiad appears to use more efficient implementations of the FLIP algorithm, both in multi-threading and memory usage.

The following chapter outlines a test scenario where inattentional blindness can be tested against the perceived quality of water when a salient object is present.

TEST METHOD

To verify the thesis theory, a comparison between videos of differing water quality with and without a salient object to attract attention needs to be tested, and the results compared.

In order to single out the influence such a salient object would have, the same video should be tested with and without the object, and with water in both a high and low quality setting.

The following 4 versions should be created; the letters following in parenthesis is the abbreviation used throughout the rest of the report to identify that setup.

1. High quality water – With object (HB)
2. High quality water – Without object (H)
3. Low quality water – With object (LB)
4. Low quality water – Without object (L)

If the salient object captivates the audience's attention they should not be drawn as much towards the water, thereby making it difficult for them to remember the quality of the water.

Following the display of each version a questionnaire should ask the test subjects to rate the quality of the water on a 7 step scale from very bad to very good. Two video examples should illustrate what should be thought of as being very bad and very good. No details should be revealed about the thesis' focus on quality and perception and neither should the comparison video be shown prior to the versions including the salient object.

The following comparison tests needs to be made:

Quality detection test (L vs. H)
High quality vs. low quality, both without salient object, would test if the test subjects are able to differentiate the quality level of the water. Each test subject needs to see both videos as the quality rating between test subjects could otherwise be subjectively influenced. The theory of change blindness could in fact influence these results, if the test subjects become blind to quality change as the water in the two videos appear similar. Consequently the test also needs to be completed in reverse order to ensure the results are not affected by the order of appearance. The average quality ratings of this test will be used for comparison in the following tests

Inattentional blindness test (LB vs. L & HB vs. H)
If the above test show that the test subject are able to differentiate the quality of the two versions, then this test should be started. If inattentional blindness does occur in this setup, then the presence of the salient object in LB and HB should captivate the audience's attention and thereby making them unsure of the quality of the water. This should consequently make the ratings more spread compared to the versions without the salient object as the test subjects become uncertain of the

quality; consequently giving a higher standard deviation compared to versions without the salient object.

An interesting observation will be to observe what happens to the quality mean in LB compared to HB, whether they approach each other or they stay separated as should be the case in L vs. H. If they stay separated the salient object either did not hold the test subjects attention enough, or they might have processed the quality of the water pre-attentively. To aid in answering that question the test subjects will be asked whether the presence of the object affected their ability to notice or remember the quality of the water. If the mean is different between the versions with and without the object, then inattentional blindness is most likely the reason. If they rate the quality of the water equal with and without the object, but answer that they had a difficult time remembering the quality of the water, the salient object most likely did not hold their attention enough. If they did not have a difficult time remembering the quality, the salient object most likely did not hold their attention at all.

Finally the test subjects will be asked to rate how critical they are towards visual effects when watching films. This question will aid in determining whether the answers given can be compared to other test subjects. A person who is used to analyzing visual effects might be more sensitive to the realism of the water simulation compared to a person who just follows the story.

DESIGN

This chapter will first outline the requirements for the product based on the analysis and the test method, and afterwards create a design for the product based on those requirements.

In order to test the thesis statement, a video presenting a high-end water simulation needs to be created. The video has to contain a salient object which the audience can attach story to, and it needs to be dramatic to retain the viewer's attention. To minimize stimuli interference it also needs to have a simple setup with as few distracting objects as possible.

As creating a complete short film with a directed story is out of scope for this project, only a scene that appears to fit in the middle of a longer film will be created. It is therefore important that the video automatically attaches story to the salient element. This way the audience should approach the video with the task to follow the object as they would follow characters and such in an ordinary film.

The setting of the video should appear as realistic as possible, without favoring specific genres such as science fiction or action. If such genres were used, the videos would also favor people with a certain mindset towards that kind of genre, which could make them overly positive or negative in their test responses.

The design requirements are as follows:

- High and low quality water simulation
- Salient object that fits in the scene
- Automatically attach story to salient object
- Simple scene setup
- Realistic setting

Based on the above list, the rest of the chapter will outline how the product should be implemented.

Overall scene decription:

To create a simple scene with only few elements and still make it dramatic, a fast flowing river was chosen as the base point. Large rocks in the water would add drama by making the water move turbulent and appear dangerous. Adding a boat as the salient object fits in the scene context and by itself add story to the video; encouraging questions from the test subjects such as "will the boat crash into the rocks?". By having the boat sail towards a waterfall the video becomes more dramatic over time, keeping the audience's attention focused at the boat. To keep the setup simple to avoid distracting objects, irrelevant background elements should be avoided, such as a sun.

Scene breakdown:

In the following, each element of the scene will be described:

The water should resemble a wild river by being splashy and having many small droplets to give it detail. In the lower detailed versions, the amount of small details can be reduced; however the main body of water should move similar in all versions. It should be rendered as close to realism as possible and use the same settings for all versions.

The boat should be simple and made of wood, this way it appears more fragile compared to a metal boat. It should move naturally down the river and react to collisions on rocks and similar.

The riverbed should appear plain without any vegetation that could distract the viewer. A matte sand-like material should form the riverbed to support the lack of vegetation. The rocks should be positioned so the boat gets close but without crashing. The rocks should have some reflection to appear wet; however highlights should be avoided as bright sports tend to attract attention.

The scene should be of short-medium length, allowing the audience to grasp the action in the scene, while not being too long, making them loose attention on the boat. Dramatic scenes in films often have fast camera cuts; this will be avoided as the camera could otherwise become a variable in the test. A scene length of 5-10 seconds is the goal.

The camera should keep the boat around its center, following it towards the waterfall, and stop shortly before the boat reaches the waterfall. Figure 8 illustrates the scene setup and the camera movement.

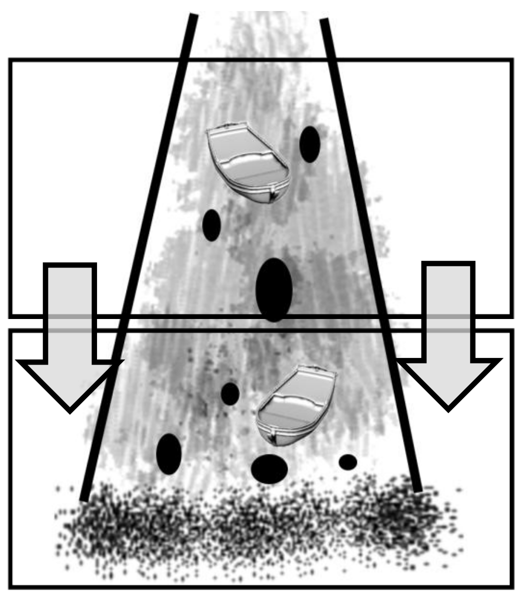

FIGURE 8 – DESIGN SKETCH WITH CAMERA FRAMING KEEPING THE BOAT IN APPROXIMATE CENTER (IMAGE BY THE AUTHOR)

To identify how to differentiate between the low and high quality version of the videos, a test was created to pinpoint which parameters were the most important according to detail level compared to simulation time in the water simulation.

To identify the impact on performance vs. simulation quality a range of tests were completed investigating the effect of the following key performance parameters; solver precision, master cell size and max time-step count. The focus was the different parameters' influence on simulation time as well as their impact on the resulting simulation quality. A simple scene with a boat sailing through a pool of water was used for all tests.

Each parameter was run through a range of iterations isolating their effect. The other parameters were left at their default value, except master cell size which was lowered from 0.1 to 0.4 to increase the test iteration speed. All tests were carried out on an Intel core i7-2600K @ 4.8ghz CPU, using 8GB of memory and a 2tb Samsung F4 5400rpm hard drive. Using the 64bit Linpack benchmarking algorithm, the CPU scored 130gigaflops in a 64bit Ubuntu 11.04 environment.

If any of the tests exceeded the 8GB of memory and started swapping to the hard drive, the test would be cancelled, this however did not happen. Based on the results, which can be found in Appendix A, the following could be concluded on the 3 parameters:

Solver precision: This parameter tells the dynamics solver how precise a solve is wanted. Increasing this value makes the dynamics solver complete more iterations in each time step, the maximum number was kept at the default at 1000. The lowest setting compared against the highest setting only showed a 16% speed difference, comparing the image sequences showed that the higher precision settings created more uniform and accurate splashes, where the lower precision showed splashes moving too fast or moving abnormally.

Master cell size: The master cell size defines the resolution of the Eularian grid, increasing the grid resolution consequently increases surface detail through particle count. Lowering this value showed an exponential increase in particle count while the simulation time only increased slowly. As the particle count increased the splashes changed radically, becoming bigger and taller.

Max time-step count: A few tests were created with different max time-steps; the solver automatically chose how many time-steps were necessary. The results showed that as the master cell size increased, more time steps were requested. By comparing images with the same master cell size but with a fixed predefined time-step count the general movement of the water body seemed similar, only by studying the still frames side by side differences appeared. With less time-steps there was less detail in the waves, and the splashes moved faster and traveled slightly longer. With double the timesteps, both the computation time and file sizes doubled as well.

These results are only comparable for this scene, but clearly show that master cell size had the biggest influence on the quality of the simulation compared to time spent. The other settings should be estimated on a per scene basis, and kept there.

IMPLEMENTATION

This chapter starts by making a short introduction to the software Naiad which was used for creating the product, hereafter the simulation graph created inside Naiad is broken down into segments and explained, to show how the water simulation was created. Finally a short view at the rendering process will be presented.

UNDERSTANDING NAIAD

Naiad is based on a client/server model and consists of two parts. The server part is the dynamics server called Naiad where all calculations occur. It's most important task is to compute the motion of bodies over time, which it uses the dynamics solver to do. A body in this context describes objects such as fluids bodies, rigid bodies, hand-animated collision bodies etc. For the dynamics solver to move its bodies it needs to complete two stages in every time-step, the transport stage and the dynamics stage. In the dynamics stage new velocities are computed according to a set of user defined parameters. At its core the dynamics solver uses an adaptive tile-based Eularian grid, which continuously changes shape in order to best surround the bodies in the scene, as seen in Figure 9.

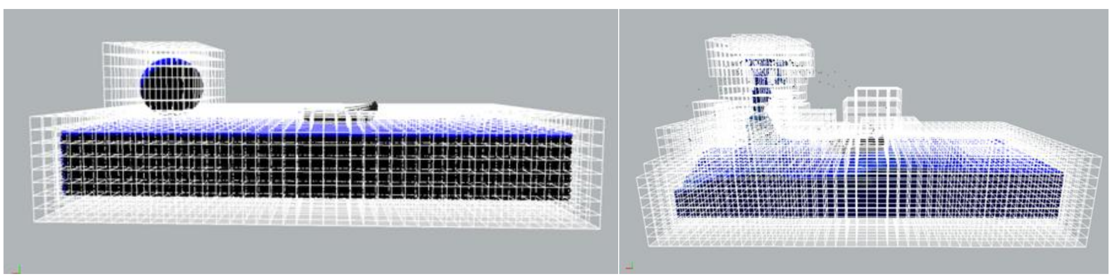

FIGURE 9 – EXAMPLE OF THE ADAPTIVE TILE LAYOUT IN NAIAD (IMAGES BY THE AUTHOR)

Being based on a Eularian grid means that data, such as velocity, is stored as a volumetric representation in fields of cells surrounding the particles, instead of residing with the particles themselves. In the transport stage all the velocity data from the fields are transferred to the particles, which subsequently move accordingly.

The client part is a graphical framework for creating and managing bodies, Open Naiad Studio, also by Exotic Matter, is such a client program. A naiad simulation is expressed in a nodal graph inside of Open Naiad Studio, where nodes representing different tasks can be wired together. The graph itself is called the "time step graph" as it portrays what happens at each time step of the simulation. The graph itself is stored in the dynamics server, the client then sends requests to the server to display, change or add to the graph. This way the graph displayed in Open Naiad Studio is always perfectly synchronized with the server.

This part breaks down how the product was created, the Naiad part is divided into 3 parts; water emission, dynamics calculations and meshing, each of these will be presented with examples images from Open Naiad Studio. A lot of nodes were joined to create the final simulation, this chapter will only talk about the most important ones.

Water Emission:

Figure 10 (left) shows the first part of the nodal graph, the water node creates a fluid, within this node all the physical parameters of the water was set, such as viscosity and density. The top emit node creates a box mesh, the second makes a mesh to levelset conversion which creates a volume representation of the surface. This volume representation uses a distance field to describe the distance from all cells to the surface. 8 of these operations were created, 1 for each of the boxes shown in Figure 10 (right). The third emit node exports the distance field, leaving all other data behind, the Dist_union01 merges two distance fields and passes it into a water emission node where particles are emitted. Finally the vel_emit adds velocity to the particles and pushes them down the riverbed. This means that water particles will be emitted from the 8 boxes shown in Figure 10 (right) with the speed and direction defined in vel_emit to make sure it flows evenly distributed down the riverbed. The 8 boxes are angled differently to make the water more turbulent.

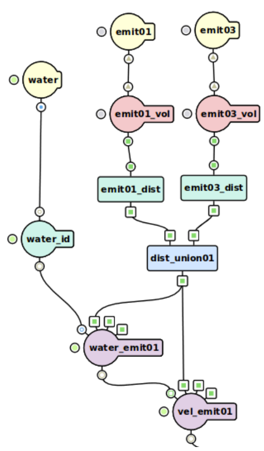

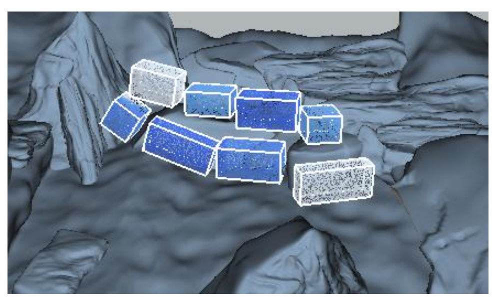

FIGURE 10 - WATER EMISSION GRAPH (LEFT) & WATER EMISSION SOURCE (RIGHT)
(IMAGES BY THE AUTHOR)

The graph above continues down the left side of Figure 11, and then again in the right side. Here the emitted particles are added the following attributes; swirl, drag, gravity and friction, these are the major parameters that affect the general behavior of the fluid. Before the fluid is passed into the dynamics node, four other branches merge with the main stream to add objects for colisions. The first imports the mesh of a boat and converts it to a rigid body, the node named boat contains all the parameters for the boat such as density and initial velocity, this controls how the boat moves through the water. One important parameter “inertial tensor” proved important for keeping the simulation of the boat stable, it lowered the impact colisions had on the rotation of the boat, this needed to be increased as the amount of particles lowered otherwise the boat would suddenly rotate abnormally and finally fly up in the air. The next branch creates a bound that deletes all particles outside of it, this ensured that particles that travelled too far from the area of interest did not consume computation power. The third branch imports the riverbed as a volume, it is nessesary that it is a volume and not just a mesh as colisions otherwise cannot occur. The last branch imports the rocks as a volume for colisions as well. Being imported as a volume means that they already went through the mesh to levelset conversion mentioned ealier.

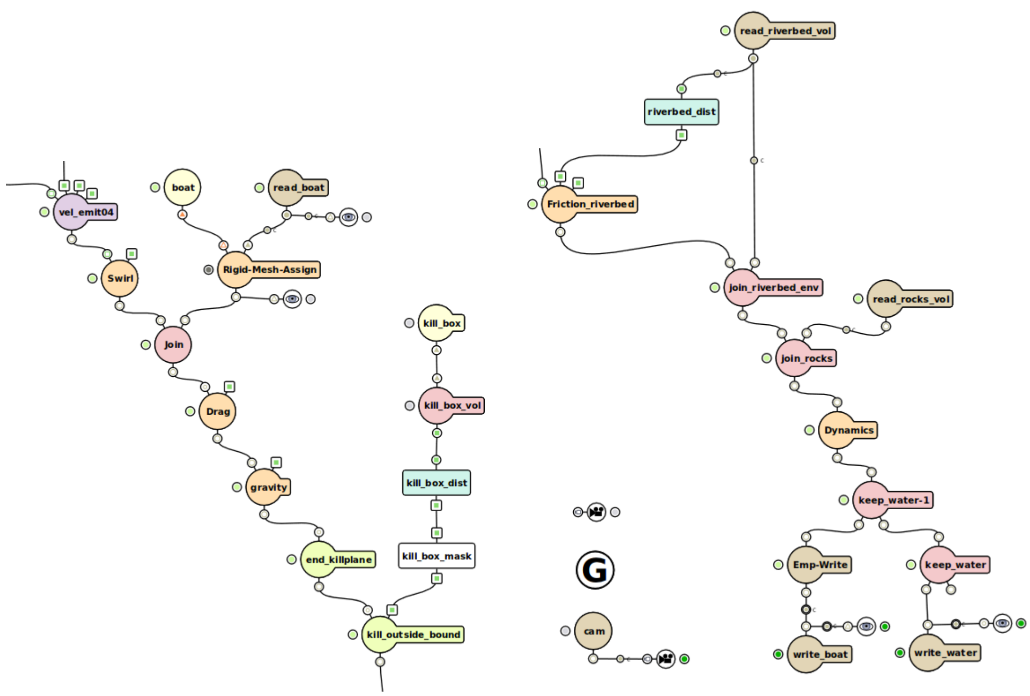

FIGURE 11 - SIMULATION AFFECTORS, SOLVE AND SAVE (IMAGES BY THE AUTHOR)

Next everything is passed into the dynamics node where the flow of the different bodies are calculated and the fluid and rigid bodies subsequently moved, this is also the node where the solver precision parameter is set. Finally the boat and the fluid is written to disk as two separate files for each frame in the simulation. The G node in the center contains the global settings, this is where the previously mentioned settings; master cell size and max time-step count are set. For each version the master cell size was the only setting that was changed, except the inertial tensor setting on the boat mentioned above. The low quality versions used a cell size of 0.4 while the high quality versions used 0.2, Table 1 shows the particle count for each of the versions at the end of this chapter.

To allow the fluid to be rendered it needed to be converted into a mesh, a range of settings allowed for custumizations of this process. Five different combinations of settings were created, the one that maintained the highest detail level in th mesh was chosen and used for all versions of the video, Figure 12 shows this process and an example of the final high quality mesh, the displayed mesh has 7.7million triangles. The triangle count for each of the versions can be seen in Table 1.

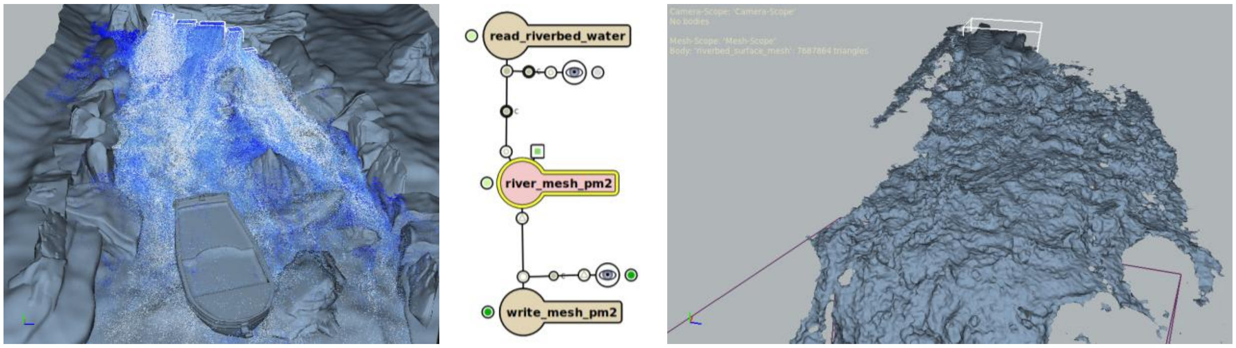

Figure 12 - The visualization of Figure 10 & Figure 11 (left)
Particle->Mesh graph (middle)
High res Mesh example (right) (images by the author)

Naiad does not contain a renderer, but the files can be converted from Naiad's propriatary format emp (exotic matter pack) to a variety of formats to allow 3rd party rendering tools to understand the files. In this case they were converted to bgeo for rendering in Mantra through SideFX's software Houdini.

The lighting setup in Houdini was kept very simple using only a high dynamic range (hdr) environment map as seen in Figure 13.

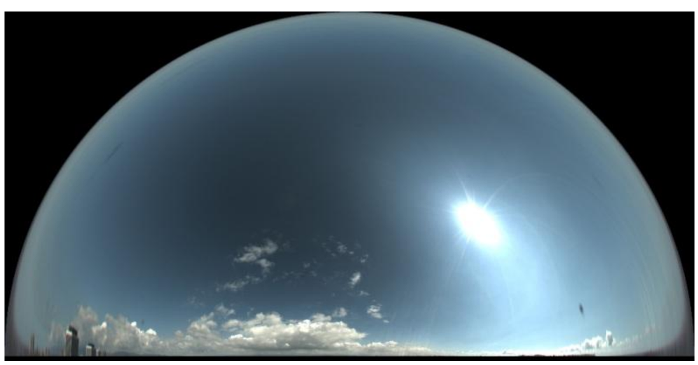

FIGURE 13 - THE ENVIRONMENT MAP USED FOR LIGHTING AND REFLECTIONS (STUMPFEL ET AL., 2011)

The rocks, riverbed and boat all used the same decal material in Houdini with low diffuse and specular reflection values, with the only exception of the rocks having a slightly higher specular value to make them shine a little as if they were wet. Each object was textures using the images shown in Figure 14.

FIGURE 14 - ROCK TEXTURE (LEFT) - SAND TEXTURE (CENTER) – BOAT TEXTURE (RIGHT) (HYPERFOCAL DESIGN,)

The water used a custom surface model to calculate physical based Fresnel refractions. It was created by Peter Quint for his tutorial on pouring water into a glass (Quint, 2011). It took the refraction index for air at 1 and water at 1.33 as inputs.

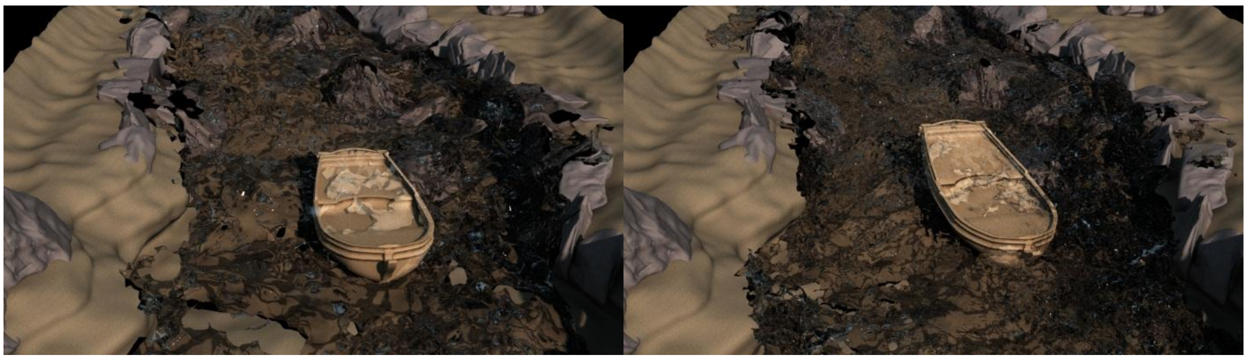

Figure 15 – Quality of LB compared to HB (images by the author)

Figure 15 compares the final rendered results of LB and HB, as it is difficult to see the difference on print, the reader is referred to the video "Compare-LB_vs_HB.mp4" on the accompanying disc. Table 1 shows the difference in simulation time as well as other data between the two versions with the boat. The difference between the versions with and without the boat was very small, thus those without the boat have not been included in the table for clarity reasons. The important value to compare is the difference in simulation and meshing time; this is how long it takes to see the result of changing some parameter. The high quality version took 174% longer to simulate and 441% longer to mesh, in this scenario this equals 24 minutes for simulation and 23 for meshing. These values may not seem like much, however many iteration needed to be completed in order to end up at the desired end result. The exact number of iteration it took to create this simulation is unknown, many iterations was halted halfway when the result of some parameter change became visible. But in total somewhere around 25 complete iterations for simulations and 5 for meshing is probably close.

The master cell size for both high quality versions was set to 0.2 for the simulation and the meshing, the low quality versions was set at 0.4, this gave an increase in particle count at 632% and an increase in triangle count of 397%. The max time step was forced to 1 as the results looked quite decent and raising it to 2 would double the simulation time for only a smaller visual difference. No changes was made in the renderer, but the difference in meshing detail made it take longer to calculate light bounces and to load in cached mesh files.

	Low quality	High quality	Increase
Particle count (202)	176314	1292363	632%
Mesh triangles (202)	1173820	5830376	397%
Simulation time (1-202)	00:13:47	00:37:47	174%
Meshing time (58-202)	00:05:15	00:28:26	442%
Render time water (58-202)	04:37:43	05:50:38	26%
Render time boat (58-202)	01:05:58	02:40:11	143%
Render time environment (58-202)	02:02:56	02:02:56	0%
Total time	**08:02:39**	**11:39:58**	**45%**

TABLE 1 - SIMULATION AND RENDERING DATA – THE NUMBERS IN THE PARENTHESIS SHOW THE FRAME RANGE OR THE SPECIFIC FRAME THE DATA HAS BEEN GATHERED FROM.

One last step was made which did not make it into the tested versions, a post splash simulation which imported the main fluid body and emitted splash particles directly from the fluid when this fulfilled a set of parameters. Examples of those parameters would be the sudden change of direction as the water collided into a rock, or the boat. At the time of testing, the movement of the splashes seemed unnaturally, so to ensure they didn't distract the audience unnecessary, they were removed. Figure 16 compares the final high quality render, with and without the splash pass. To more clearly see the difference the reader is referred to the comparison video "Compare-HB-Splash_vs_HB-NoSplash.mp4" available on the acompanying disc. The splash particles were rendered as simple points.

FIGURE 16 - FLUID WITH SPLASH VS FLUID WITHOUT SPLASH (IMAGES BY THE AUTHOR)

This product is the second iteration created for this project; Appendix D: First product iteration shows the first product iteration where a meteor was meant to make impact with the ocean and the following waves would then hit a boat in the water. As the theory of inattentional blindness was studied further the product was however redesigned as the impact would attract too much attention away from the boat. A video named "product - 1st iteration.mp4" of the first iteration can be found on the accompanying disc.

TEST

Legend:	
L:	Low quality without boat
H:	High quality without boat
LB:	Low quality with boat
HB:	High quality with boat

The test method stated that two comparison tests should be made to properly isolate and test inattentional blindness; Quality detection test (L vs. H) and Inattentional blindness test (LB vs. L & HB vs. H).

The quality detection test should validate that the test subjects were able to differentiate the quality between the different rendered versions of the water. It was important that the same test subject rated both the L and H version, as rating quality is very subjective, so to ensure that the order in which the videos was shown did not influence the ratings, half the test subjects were shown the videos in reverse order. The quality ratings found in this test should then be used for comparison in the inattentional blindness tests, which compares if the test subjects rate the versions with the boat differently than the versions without.

The purpose of the tests was to identify tendencies and prove the concept of the idea and the test; consequently the sample size for each comparison was set relatively low at 30.

Instead of completing 3 tests with 30 test subjects for each; namely L vs. H, LB vs. L and HB vs. H, the different comparisons were gathered in the groups outlined in the table below to reduce the total number of test subjects needed. This was only to reduce the load on testing.
The videos testing inattentional blindness had to be tested on test subjects who had no idea that the test worked with quality and perception, consequently they had to be shown first.

	Video order	Samples
Group 1	LB -> L -> H	15
Group 2	LB -> H -> L	15
Group 3	HB	30

Table 2 - Test groups

The target group used for the testing was male and female in the age range of 18 and 34, as those have been defined to be the largest movie going segment (Pigott, 2011). Furthermore people with knowledge of fluid simulations were avoided to avoid biasing the results.

TESTING PROCEDURE

This part presents a step by step example of how the test progressed from start to finish.

1. Test subject walked in and was sat down and told to pay attention to a short video.
2. The first video was started with a 5 second film leader.
3. As soon as the video finished they were handed a questionnaire and explained the first question to rate the quality of the water they just saw on a scale from 1 to 7 – Another monitor was turned with two repeating video streams representing what should be considered 1 and 7 on the scale. This monitor kept cycling the comparison video throughout the test. However angled so it was not visible as the other test videos ran.
4. When questions related to the first video had been answered, the second video was started.
5. Upon completion, they turned the page and answered all questions related to video 2.
6. Step 4 and 5 repeated for video 3 where they furthermore completed the rest of the questionnaire.

The repeating comparison video can be seen at the accompanying disc named "Compare-VLB-HB.mp4". A shortened version of the questionnaire is shown in Appendix B: Questionnaire summary.

TEST RESULTS AND ANALYSIS

Quality detection test (L vs. H)

The purpose of this test was to check if the test subjects were able to distinguish the quality of the low and the high quality water. The test was split into two groups in order to verify if the order of appearance had any influence on the test subjects' responses. The ratings for low quality in both groups were largely similar as seen in Table 3. The standard deviation was slightly above 1, this however was expected due to the subjective rating of quality as also mentioned above.

Quality detection test	Mean	St.Dev.	Samples
Low quality water (L)			
Group 1	4.33	1.19	15
Group 2	4.07	1.06	15
Combined	**4,20**	**1,14**	**30**
High quality water (H)			
Group 1	6.13	1.02	15
Group 2	5.92	0.76	12
Combined	**6,04**	**0.92**	**27**

Table 3 - Test results for the Quality detection test

In the high quality comparison, Group 1 rated the quality at 6.13 with 1.02 in standard deviation, slightly higher than group 2. However in group 2, three people rated the high quality at 3 and blamed shading artifacts from the rendering, with a sample rate of only 15 per group this had a high influence on the results, giving a mean of 5.3 and a standard deviation of 1.35. As the shading artifacts were also present in the two videos defining the scale, these three responses were ignored. This made the results look similar to those from Group 1; the results displayed in Table 3 confirm this. Figure 17 shows a histogram for both the low and high quality version, each comparing the responses from both groups.

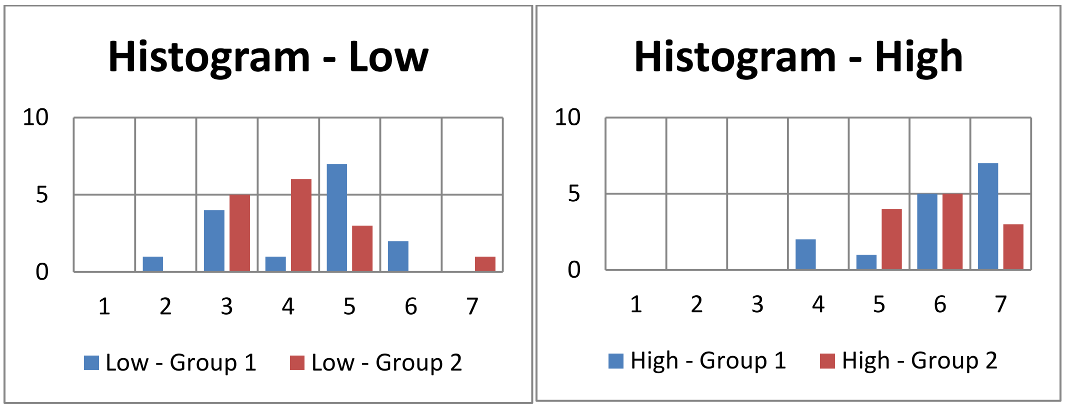

FIGURE 17 - HISTOGRAMS COMPARING THE RESPONSES FROM THE TWO GROUPS IN BOTH THE LOW AND HIGH QUALITY VERSIONS OF THE QUALITY DETECTION TEST

The above histograms and table show that there were differences in the responses for both low and high quality. However the results seem close enough to indicate that showing two consecutive videos to the same test subject did not influence the quality rating of each one. Merging the results from the two groups show an average quality difference between the low and high quality of 1,83, which clearly show that the test subjects were able to differentiate the quality between the two. Figure 18 below shows the results combined from the two groups. As can be seen the low quality responses lean towards the left, whereas the high quality responses lean towards the right.

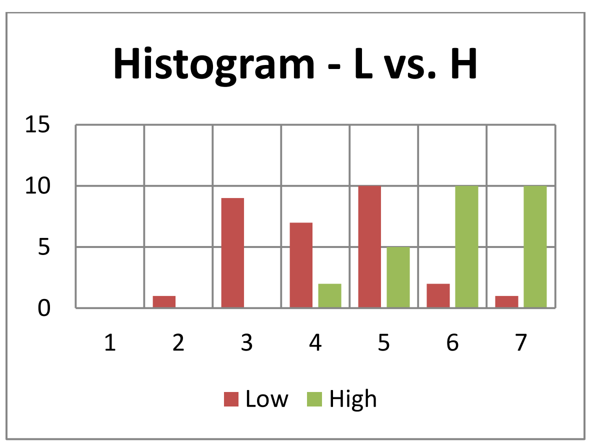

FIGURE 18 - COMPARISON BETWEEN LOW AND HIGH QUALITY USING THE COMBINED RESPONSES FROM GROUP 1 AND 2

The mentioned shading artifacts, exemplified in Figure 19, were present in all versions and were a result of a layering error in the rendering. The black shape is water behind the rock which was incorrectly composited on top of the rock. The reason for the attempt to layer the rendering was two things, first to be able to optimize the rendering settings to each object and second to be able to reuse the riverbed and the rocks in all 4 versions to save rendering time.

FIGURE 19 - EXAMPLE OF SHADING ARTIFACT (IMAGES BY THE AUTHOR)

Inattentional blindness test (LB vs. L & HB vs. H)

The mean and standard deviation results of L were very similar to those of LB as seen in Table 4, so at first glance it would seem that the theory of inattentional blindness was ineffective in this setup. In the test method it was hypothesized that if inattentional blindness was effective, then the standard deviation would be bigger in the version with the boat, this was not the case. When studying the responses however, the majority of the test subjects responded yes (14/30) or might have (8/30) to the question "*Did the presence of the boat affect your ability to remember the quality of the water?*"". This makes it seem that they were affected by it, but the boat didn't captivate their attention enough.

Inattentional blindness test	Mean	St.Dev.	Samples
Low quality water (L)	4,33	1,19	30
Low quality water with boat (LB)	4.23	1.05	30

Table 4 - Test results for the Inattentional blindness test

Another explanation would be that the water was just too close to the boat; it splashes against it and moves inside the boat. 11 of the 30 test subjects commented that they looked at the water inside the boat, some even thought the boat was submerged in water. Either way it grabbed their attention and by doing so led their attention straight to the water, which obviously defeated the purpose. Appendix C: Test shows additional examples of comments from the test subjects.

After this test a dilemma had to be solved, either to continue with testing the HB version, or to redesign both LB and HB for a later test where both the shading artifacts were removed, and the behavior of the water inside the boat was fixed.

As the purpose was to create a testing scenario that encouraged inattention blindness, continuing with a comparison against HB would likely just provide invalid results, if the boat indeed directs the attention directly to the water as it would seem from the comments. The comparison would in that case largely be based on the perceived quality of the water inside the boat. The water in the HB version seemed to dissipate more realistic due to a higher particle count and meshing resolution. In the LB version the low meshing quality made the water blobs bigger, combined with fewer particles the water appeared to dissipate faster. The accompanying disc contains a video named "Compare-LB_vs_HB.mp4" which shows LB and HB side by side to illustrate the difference.

Based on the above discussion, it was decided that the best way to proceed was to make a second iteration of the product where the dissipation issue of the water in LB was fixed as well as the shading artifacts present in both videos. This or other possible continuations of the current findings were however left to future studies. The next chapter discusses the currents results as well as the general project process. Following this is a future perspectives chapter that talks about other ways this project can be continued by the author or other people in the future.

DISCUSSION AND CONCLUSSION

The motivation for this project emerged because of the immense computational requirement of high-end water simulations for motion pictures. Especially the question of whether the high quality level set by industry professionals was too excessive compared to the quality perceived by the regular audience. The quality detection test showed that the test subjects were able to differentiate the quality between the low and high quality version created for this project; however the rated quality increase was much lower than the increase in simulation and meshing time. This alone indicates how the increased computational efforts are not always appreciated by the audience. In this project the simulation time alone was a mere 38 minutes for the high resolution; however in high end motion pictures the simulation quality is often set considerably higher. Exotic Matter did a case study of the motion picture Shark Night 3D from 2011, directed by David R. Ellis (Ellis, 2011), where Igor Zanic, the lead Naiad TD, said that the most complex shot in the film took over a week of continues simulation to complete (Exotic Matter, 2011a).

The overall goal was however not to investigate how an audience rate the quality of water, but to see if they even notice the quality when an object more important to the story was present. There was no directed story is the videos created, but based on investigations into visual perception it was hypothesized that certain categories of objects automatically attract more attention than other. The hypothesis essentially said humans over animals over objects over environment. During the testing the majority of test subjects stated that the presence of the boat had, or might have had, an influence on their ability to recall the quality of the water. However despite this, the difference in mean quality rating between the video with the boat and the video without was very small.

Some test subjects mentioned “shading artifacts” as a distraction in the video, what was referred to was rendering errors due to an incorrect layering setup which caused very dark water from behind the rocks to appear on top of the rocks. Some people mentioned these directly, but others were most likely indirectly affected as well, as it in retrospect did attract unnecessary attention. Furthermore several test subjects mentioned that they looked at the water inside the boat, which as well had was affected by an incorrect layering operation. Figure 20 compares an image with the incorrect layering setup used in the tests, with one where the entire scene was rendered in one layer by the same render node. It is easy to see the difference, the HB version has black spots several places across the scene and furthermore the water inside boat looks different.

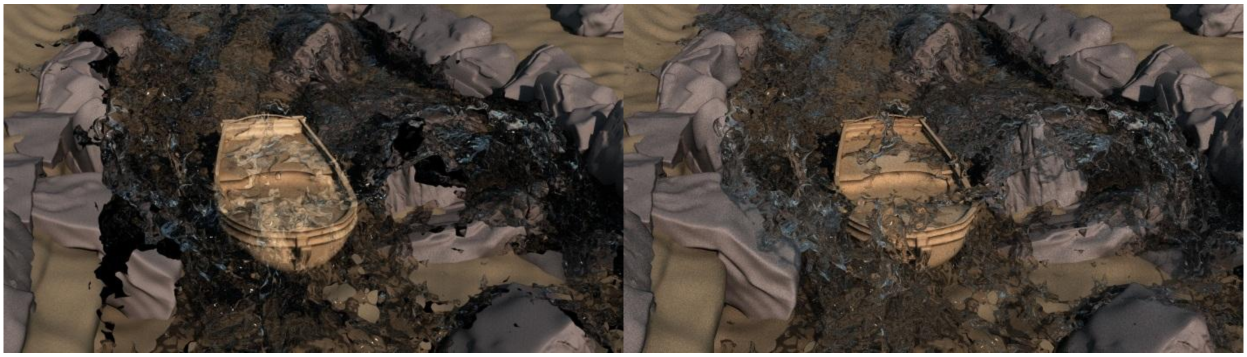

FIGURE 20 – HOW THE RENDER LOOKED (LEFT) COMPARED TO HOW IT SHOULD HAVE LOOKED (RIGHT) (IMAGES BY THE AUTHOR)

After the first inattentional blindness test it was apparent that the shading artifacts as well as the water inside the boat, both were objects that attracted too much unnecessary attention, consequently both needed to be fixed before a proper inattentional blindness test could be completed. The shading artifacts could be fixed by either rendering everything in one layer, or by making sure that no detail gets lost in the layering process. The latter way would be most optimal to allow different rendering settings for each object, such as sample rate.

The problem with the water inside the boat was two things, first an incorrect compositing operation was used which made the water on the boat appear different than the rest. Second the water dissipated abnormally fast in the low quality version. The first was easily fixed by changing the compositing operation from “screen” to “over”. The dissipation problem was the product of a low simulation and meshing resolution, this could be concealed using splash particles. Those were removed before testing as they behaved unrealistic, but a new simulation could fix that.

The second problem with the boat was the boat itself. The boat did attract the attention of the audience as intended, but by doing so led them straight to the water inside the boat. If the unrealistic appearance was to be fixed the audience would still be looking straight at the water, in which case an even more salient object should most likely be appearing on the boat to make sure the audience looked at that, instead of the water.

Inattentional blindness is an area which has received increased attention from the computer graphics industry, especially in rendering. This project have presented theories that using inattentional blindness to allocate computational time should not be limited to rendering, but might in fact be applicable in all areas of the computer graphics production of a shot.

The thesis statement for this project stated:

> **To what degree is the perceived quality of a water simulation affected by the presence of objects related to the story in a video.**

Although the quality ratings did not show a difference in perceived quality, the comments showed the majority of test subjects were distracted by the presence of the boat when trying to remember the quality. It can consequently be concluded that the presence of objects with higher saliency does affect the quality perception of other objects, such as a water simulation. However to determine the degree to which the perceived quality is affected, more investigations need to be completed.

Continuing investigations in this area could potentially lead to a framework usable in the visual effects industry to target which computer graphics elements in a scene receive the most attention from the audience and should consequently receive the most attention during production.

The following future perspectives chapter presents examples of how this, or other projects, can proceed from the findings of this thesis.

FUTURE DEVELOPMENT

From where this project ends, there are several directions other projects can proceed.

The current test can be finished by improving the already mentioned areas of the videos. Furthermore the saliency of the boat can be increased by adding some sort of action on it, to further ensure that the test subjects do not look at the water on the boat.

Another approach to solving the thesis statement could also be made. This project started by creating one product but ended with using another because it was anticipated that the water in the first iteration would attract too much attention. The design used in the project have a boat sailing down a river, this directs the audience's attention to a place very close to the water. It would be interesting to see what influence the proximity between the water and the salient object has on the perceived quality. The scene could be changed to show a river flowing underneath a bridge with a camera focusing on a conversation between two people on top of the bridge. This different setup would push the activation of the water further back.

Doing a range of tests on different proximities, qualities, scene complexities etc. could then be gathered in a framework usable for visual effects companies to distribute production time based on the findings in this framework.

WORKS CITED

Bordwell, D. (2011). *The eye's mind.* Retrieved from http://www.davidbordwell.net/blog/2011/02/06/the-eyes-mind/

Brackbill, J. U., Kothe, D. B., & Ruppel, H. M. (1988). Flip: A low-dissipation, particle-in-cell method for fluid flow.

Bridson, R. (2008). *Fluid simulation for computer graphics.* Wellesley, Massachusetts: A.K. Peters, Ltd.

Bridson, R., & Yongning, Z. (2005). Animating sand as a fluid. *ACM Trans. Graph., 24*(3)

Cater, K., Chalmers, A., & Ward, G. (2003). Detail to attention: Exploiting visual tasks for selective rendering. *Proceedings of the 14th Eurographics Workshop on Rendering,* Leuven, Belgium. 270-280.

Cater, K., Chalmers, A., & Ledda, P. (2002). Selective quality rendering by exploiting human inattentional blindness: Looking but not seeing. *Proceedings of the ACM Symposium on Virtual Reality Software and Technology,* Hong Kong, China. 17-24. doi:http://doi.acm.org/10.1145/585740.585744

Briggs, C., Fleiss, M., & Howell, L. (Producers), & Ellis, D. (Director). (2011). *Shark night 3D.* [Motion Picture] United States: Rogue.

Exotic Matter. (2011a). *Case study: Shark night 3D.* Retrieved September 30, 2011, from http://www.exoticmatter.com/shark-night-3d/

Exotic Matter. (2011b). *Home.* Retrieved 12 May 2011, 2011, from http://www.exoticmatter.com/

Failes, I. (2011a). *Framestore's heavens and horcruxes.* Retrieved from http://www.fxguide.com/featured/framestores-heavens-and-horcruxes/

Failes, I. (2011b). *Your highness: A little bit of framestore magic.* Retrieved, 2011, from http://www.fxguide.com/featured/your-highness-a-little-bit-of-framestore-magic/

Healey, C., & Enns, J. (2011). Attention and visual memory in visualization and computer graphics. *IEEE Transactions on Visualization and Computer Graphics,* doi:10.1109/TVCG.2011.127

Hyperfocal Design.*Sample textures.* Retrieved September 14, 2011, from http://www.hyperfocaldesign.com

J., D., & W., H. G. (1989). *Visual search and stimulus similarity.* Washington, DC, ETATS-UNIS: American Psychological Association.

Lait, J. (2010). *Introduction to FLIP fluids.* Retrieved October 2, 2011, 2011, from http://www.sidefx.com/index.php?option=com_content&task=view&id=1834&Itemid=132

Liu, G. R., & Liu, M. B. (2003). *Smoothed particle hydrodynamics.* Singapore: World Scientific Publishing Co. Pte. Ltd.

Musker, J., & Clements, R. (Producers), & Musker, J. and Clements, R. (Directors). (1992). *Aladdin.* [Motion Picture] United States: Walt Disney Pictures.

Next Limit Technologies. (2010). *HYBRIDO - white paper*

Pigott, M. (2011). *Understanding your target audience.* Retrieved, 26, from canadianfilmmaker.com

Priscaro, M. (2009). *NVIDIA quadro and "2012" change our view of the world as we know it.* Retrieved from http://blogs.nvidia.com/2009/12/nvidia-quadro-and-2012-change-our-view-of-the-world-as-we-know-it/

Quint, P. (2011). *Basics : Glass / ice.* Retrieved Septeber 14, 2011, from https://sites.google.com/site/pqhoudinitutorial/

SideFX Software. (2011). *What's new in houdini 11: Dynamics and effects.* Retrieved from http://www.sidefx.com/docs/houdini11.0/news/11/effects

Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception, 28*(9), 1059-1074.

Simons, D. (2000). Current approaches to change blindness. *Visual Cognition, 7*(1), 1-15.

Stumpfel, J., Jones, A., Wenger, A., Tchou, C., Hawkins, T. & Debevec, P. (2011). *Direct HDR capture of the sun and sky.* Retrieved September 29, 2011, from http://gl.ict.usc.edu/skyprobes/

Sundstedt, V., Debattista, K., Longhurst, P., Chalmers, A., & Troscianko, T. (2005). *Visual attention for efficient high-fidelity graphics*. Budmerice, Slovakia: ACM.

Transport for London. (2009). *The awareness test.* Retrieved from http://www.awarenesstest.co.uk/

Treisman, A. (1985). Preattentive processing in vision. *Computer Vision, Graphics, and Image Processing, 31*(2), 156. doi:DOI: 10.1016/S0734-189X(85)80004-9"

Verrier, R. (2011). *Fade out for visual effects.* Retrieved, 24, from http://articles.latimes.com/2011/feb/01/business/la-fi-ct-visual-effects-20110201

Yarbus, A. L. (1967). *Eye movements and vision.* New York:

APPENDIX

APPENDIX A: PARAMETER TEST DATA

MCS	Sim start	Sim end	Total sim time	Total part. count	Max timestep/frame	Total Timesteps	Total file size
1	15:24:23	15:24:47	0:00:24	35763	5	48	79MB
0,9	15:24:47	15:25:16	0:00:29	44862	5	49	102MB
0,8	15:25:16	15:25:50	0:00:34	58121	5	48	128MB
0,7	15:25:50	15:26:38	0:00:48	77432	5	54	191MB
0,6	15:26:38	15:27:51	0:01:13	108227	5	64	300MB
0,5	15:27:52	15:29:38	0:01:46	156094	5	77	511MB
0,4	15:29:38	15:32:58	0:03:20	221472	5	107	997MB
0,3	15:32:58	15:41:36	0:08:38	401775	5	143	2.5GB
0,2	18:31:38	18:55:38	0:24:00	872335	5	158	6.4GB
0,2	6:41:55	7:07:55	0:26:00	872386	10	171	7.0GB
0,15	1:59:35	3:04:10	1:04:35	1749640	5	197	16.1GB
0,15	11:14:34	12:21:11	1:06:37	1749761	8	202	16.5GB
0,15	3:04:11	4:30:01	1:25:50	1749500	20	258	21.3GB
0,1	2:29:29	6:41:54	4:12:25	5047199	5	224	50.5GB
0,1	7:09:20	15:16:41	8:07:21	5048705	10	369	84.9GB
Dynamics Precision							
0,1	15:41:36	15:44:28	0:02:52	221400	5	96	894MB
0,2	15:44:29	15:47:34	0:03:05	221400	5	103	957MB
0,3	15:47:34	15:50:40	0:03:06	221400	5	104	967MB
0,4	15:50:40	15:54:05	0:03:25	221400	5	111	1010MB
0,5	15:54:06	15:57:20	0:03:14	221400	5	105	976MB
0,6	15:57:21	16:00:31	0:03:10	221400	5	103	956MB
0,7	15:29:38	15:32:58	0:03:20	221400	5	107	997MB
0,8	16:00:32	16:04:05	0:03:33	221400	5	109	990MB
0,9	16:04:05	16:07:28	0:03:23	221400	5	105	974MB
1	18:26:41	18:30:05	0:03:24	221400	5	106	988MB

Part 1: Asked after every video shown

Question	1 – **Very Bad**	2	3	4 – Medium	5	6	7 – **Very Good**	Don't know
How would you rate the quality of the water on a scale defined by the two example videos shown?								

Additional comments (Optional)

Part 2: Asked after video 2 and 3

Did you notice any difference in the quality of the water between the first and the second video?

	First video had better quality
	Second video had better quality
	Seemed to be the same
	Don't know

Additional comments (Optional)

Part 3: Asked in the end of the test

Did the presence of the boat affect your ability to remember the quality of the water?

	Yes
	It might have
	No
	Don't Know

Please elaborate your answer:

Age _______ Gender _______

Rate how critical you are towards the realism of visual effects in movies.

(Such as explosions, water floods, collapsing buildings etc.)

	Not critical	Ex. - It's not something I notice at all, I just follow the story.
	Slightly critical	Ex. - Sometimes I notice it, but it generally does not distract me.
	Somewhat critical	Ex. - I notice visual effects which are out of place, sometimes it distracts me.
	Critical	Ex. - I often notice visual effects, it distracts me from the story.
	Very critical	Ex. - I actively search for visual effects and think about how they were created.

Additional comments on anything related to this test: (Optional)

APPENDIX C: TEST COMMENTS

LB

- It looked like the boat was a bit too deep in the water, but other than that it looked pretty realistic.
- There were water inside the boat
- Seemed very chunky sort of like crystals, rather than fluids

L

- Looked like it was flowing too slow and looked too thick

H

- The shadows made the water look very dark and almost oil like.
- The water became too dark and clumpy
- The particles almost seemed like fire in this one (comment: He is talking about the shading on the edges)

Presence of boat comments

- Focused a lot on it and wondered why there was water in the boat. So a clear distraction.
- I was annoyed by the water flowing thru the boat
- I focused on the water inside the boat
- I was primarilly irritated with the fact that the boat didn't seem solid and was see thru.

All questionnaire responses are gathered in an excel file on the accompanying disc.

These 8 images are screenshots with a 24 frame interval from the product as it first appeared, the idea was to exchange the ball of water with a small asteroid, but the idea was dropped after more research in perception was made. The design would draw too much attention away from the boat and on to the impact with the water.
This sequence is 181 frames; the rest of the images as well as a video have been included on the attached disc. It has been mentioned here to illustrate the process the project has been through.