

Modelling of Turbulent Water over Natural Terrain

Nathan Holmberg and Burkhard Wünsche
Graphics Group, University of Auckland, Auckland, New Zealand
{nhol021,bwue001}@ec.auckland.ac.nz

Abstract

Water phenomena are some of the most visually spectacular effects found in nature. This paper presents a novel efficient hybrid method to model turbulent water such as fast flowing rivers and waterfalls with the intent that the model can be used as part of a larger environment or scene. The model presented uses hydrostatic theory to incorporate a 2D height field and a particle system to model respectively the main volume and spray of turbulent water. The user is able to submit any environment formed from spheres and panels making the solution very flexible and adaptable. Our results show that the model provides a nearly realistic simulation of turbulent water and for simple scenes interactive speeds are possible which compares favourably with alternative techniques.

Keywords: physically-based modelling, turbulent water, water simulation

1 Introduction

The complexity and power of water flow in nature is both impressive and beautiful. It is also a part of our everyday lives and so well known to the human perception that unrealistic motion is easily discernable.

It is no surprise therefore that much effort has been applied in computer graphics to try to capture water in a convincing way. This is far from trivial as most phenomena can be attributed to complex molecular interactions that occur trillions of times a second, and as such are beyond the modelling power of today's computers. Instead it is the goal of computer graphics to produce models with an underlying basis in physical principles that create, as accurately as possible, a large scale approximation of these local interactions.

As research in this area has evolved models have moved from only being able to represent certain effects, such as the motion of deep water waves to the exclusion of all else, to more general models with a firmer basis in hydrodynamics capable of realistic motion that follows expected behaviour in a range of situations. These more recent models allow animators to specify environments and start conditions and the model will do the rest.

The purpose of this work is to build a general model with a focus on the natural movement of rivers, rapids and waterfalls. As such turbulence, spray, and the like should be accounted for while not being treated as special cases.

While not yet being complete and fully realistic, our model has been able to produce results that show its potential.

2 Previous Work

The earliest work on modelling water in computer graphics used mathematical models and explicit

functions to model surface behaviour. Some, while obeying physical equations, were very inflexible outside their intended scope such as the modelling of deep sea waves by Schachter [1] while others tried to model surface phenomena in 2D velocity fields [2].

Later work introduced rudimentary particle systems. Initially particle interaction was ignored and particles simply bounced around the environment [3][4]. This produced reasonable effects for waterfalls and other instances where particle movement had enough energy to break the molecular level bonds as needed. In particle systems the equations for interaction can be complex and computationally expensive as each particle may be affected by every other resulting in a computational complexity of $O(n^2)$ which is unacceptably slow when dealing with hundreds of thousands of particles. However, without inter-particle forces volumes of water can not be modelled well and the usually identifiable effects of adhesion and cohesion are missing. Some of the work at modelling these effects includes Miller and Pearce's [5] connected particle system which used forces to imitate soft collisions based on the difference in particle's positions. This works well for viscous fluids, as intended by the authors, but is unsuitable for the number of particles or scale necessary to model large bodies of water.

Kass & Miller [6] introduced column systems which use simplified flow equations, based on hydrostatics, between columns and treat the water volume as a 2D height field. This implicitly allows for the modelling of surface phenomena such as waves and can efficiently model large bodies of water. O'Brien & Hodgins [7] extended this model by incorporating the interaction with external objects and splashes. However, a column's height is still represented by one height variable meaning that vertical isotropy is assumed and only very simple (flat) environments can be used. Mould & Yang [8] furthered this work by

dividing each column into a user defined number of cells, relaxing the assumption of vertical isotropy. Another improvement was made by allowing for more complex environments. Even with the use of cells, however, column systems are still only a two and a half dimensional model at best and many effects such as the curling of waves cannot be reproduced. There has also been some work at applying the 2D version of the Navier-Stokes equations to column systems [9]. The most recently developed simulation methods use the full Navier-Stokes equations to compute motion in a 3D grid of voxels, examples of which can be found in Enright et al. [10] and Foster & Fedwik [11]. The solution is very expensive but also creates the most realistic effects. It is obtained by calculating the velocity of flow out of the 6 faces of each voxel using the pressure, temperature etc. of each voxel. By calculating which voxels contain water, objects, or air, extremely realistic simulations are obtained for complex environments. This method has been used in movies such as Shrek and its difficulty is perhaps illustrated in Jeffrey Katzenberg's statement that the pouring of milk into a glass was the hardest shot in the movie [10][11]. This approach does have its limitations of course with perhaps the largest being the loss of information within cells, meaning that the grid size plays an overly important role in the accuracy of the model.

3 A Hybrid Model for Turbulent Water

Initial tests with a particle only model showed that the computation time for such a system is prohibitive for large bodies of water. Instead we chose to use different systems for spray and volume to create a faster and more believable model.

3.1 Volume Model

Continuing the work of O'Brien & Hodges and Mould & Yang [7][8] the volume model, representing the main body of water in our model, is a column based system. Using columns holds the advantages of easy surface creation as the top of all columns are known and less flow calculations are needed so the system is less computationally intensive.

Our model divides the environment into equally sized squares which form the base of columns as in figure 1. All columns start with a user defined height which then varies over time dependent on the calculated flows. Source and sink columns are the only ones to retain their heights allowing for in and out flows to the system. Pipes are then created between each of the eight adjacent columns and each cell that could overlap during the course of the simulation. Pipes are also created between the cells of one column and the air above the adjacent columns. At this point the system is ready to begin simulation.

The underlying basis for the flow equations that are used in this simulation is the science of hydrostatics,




Figure 1. Column Generation from environment. Arrows indicate position of some pipes.

which describes the pressure of fluids at rest. The equations related with this approach are simple, both to understand and to compute, and as a result are easy to implement.

For any column in the grid the hydrostatic pressure can be calculated from the equation

$$Q = h\rho g + \frac{1}{2}\rho v^2 + (p_0 + E)$$

which is based on the work of Bernoulli and where Q is the total pressure, h the height of the column, g the acceleration due to gravity, ρ the density of the fluid, v the velocity of flow, p_0 the air pressure and E the pressure arising from external forces which together form the pressure energy term. In this case the height of the column is the height above some arbitrary point in the world so long as the same point is used for all columns.

Using the pressure differences between cells it is possible to calculate the acceleration and from that the flow that should occur between cells. The final equation for the flow velocity (η) through a pipe is

$$\eta = f\eta_0 + \Delta t \left(\frac{Q_{head} - Q_{tail}}{\rho l} \right)$$

where l is the pipe length, f is a friction coefficient (as suggested in [8]), and η_0 is the flow in the previous time step. An interesting point to note is the lack of any viscosity parameter in this equation. This is because one of Bernoulli's assumptions is that the distances between points of measurement are so small that viscous losses are negligible. Instead the friction coefficient used allows energy to slowly escape from the system. While not physically justified this parameter serves as an ad-hoc method of including viscosity and we found that setting it to 0.995 gave good results in our examples.

Using the flow calculated for the pipe the volume of water that should be moved through it is calculated by:

$$V = \Delta t \eta c$$

where c is the cross-sectional area of the pipe, or the

amount of overlap between the cells. Because mass is to be conserved the volume removed from one column is the same as that added to the other. Care must also be taken not to allow a volume of less than zero to occur.

This system is very fast considering the mass of water that is being represented but there are several problems.

Columns pose a problem as a representation because one of the classical characteristics of turbulence is that it is a three dimensional feature [1]. While using the 'cells' given by Mould & Yang relaxes the assumption of vertical isotropy this goes a long way to explaining some of the model's inability to simulate certain situations. A second problem arises from the fact that turbulence is a feature of flow and not of the fluid 'at rest'. This means that while hydrostatics may be easy to use the equations generated for flow are incomplete and ignore many of the visible characteristics of water such as viscous shear stresses. This is perhaps the largest flaw in the model and something that lends itself to further research

3.2 Spray Model

The spray model is used to model water as it breaks free of the main volume of water. There is no easy physical solution to when spray should be created and as such assumptions must be made instead. Earlier work with column systems were concerned mainly with generating splashes from hitting objects and as such used vertical velocity thresholds for generating spray. Because we also want to model waterfalls our assumptions use research about the heights of wave crests before they become unstable (when the wave height is 0.78 of the water depth) [12]. This obviously allows waterfalls to form easily but also works for rapids as large flow velocities form 'spikes' of water that while erroneous are then turned into spray due to their large heights.

The spray system begins its evolution when particles are generated. First the number of particles (or volume) needs to be determined. Using the formula to calculate flow through a weir [13] it is possible to determine the required flow rate and hence volume:

$$flowrate = \frac{2}{3}BH^{\frac{3}{2}}\sqrt{2g}$$

where B is the base length, H the height and g the acceleration due to gravity. The volume to pass through in this time step is then given by $flow\ rate * time\ step$. This determines the number of particles to be created as all those generated are of a user defined volume except the last which needs to be the remainder of volume to be moved as mass must be conserved and not created. Depending on the scale and resolution of the model being used this can be set to achieve the best looking results. The position of the particles is also easily determined and is set at a random position in the face that the particle is being generated from. The final initial variable that is

needed for each particle is velocity which is generated from the flow rate equation above. In this case the velocity of flow through a column's face is found to be flow rate / face area. Flows within the column structure may mean the velocity should not be perpendicular to the front of the face; to account for this the average flow from surrounding columns is used to give a direction to the scalar velocity calculated above.

In many cases the velocity imparted to the particle should not only be horizontal but also include an initial vertical velocity. To do this the difference in total height between the column for which particles are being generated and the column behind is used in the classic formula

$$v^2 = u^2 + 2as$$

Where v is the current velocity, u is the initial velocity, a the acceleration, in this case gravity, and s the distance covered. While only an approximation this approach does manage to provide a more believable representation.

One of the methods that were initially considered to help create the illusion of water pooling and incompressibility was the use of cohesion. Intermolecular bonds are simulated by creating small forces between water particles which attract and repel neighbours in the effort of keeping an optimal distance apart. However, we found that using cohesion within the particle system is inefficient and inaccurate. Furthermore other authors suggest that cohesion is unnecessary due to the water's low viscosity so long as the movement is turbulent, such as is the case with spray [4][14]. Consequently this extension was excluded when the particle and column systems were combined and the only active force during a time step is gravity.

It remains to deal with the collision detection for particles. Collision detection with the water volume boundary is explained in the next section. It remains to handle particle collisions with columns.

After calculating which column the particle is above (or within), the particle's y-coordinate value is tested to see if it should be absorbed into the column. However, simply increasing the column's volume and destroying the particle was found not to be accurate enough as this can force the column's height up, often absorbing more particles in the process. While this is not a problem for small scale splash effects considerable problems occur when modelling, for example, waterfalls where there are many particles hitting at any one time. After trying to spread the volume of a colliding particle over several columns it was found that by instead modelling the force of impact and subsequent pressure increase in the column not only was the problem reduced but more realistic effects were generated. The equations used are the same as for external objects colliding with the water presented in Mould and Yang. The force on the object consists of two terms:

$$F = -v\mu - V\rho g$$

where v is the velocity of the object, μ is the viscosity, V the volume of water displaced, ρ the density of the fluid and g the acceleration due to gravity. The first term describes the force of the fluid on the particle and the second is the force due to buoyancy. Because the forces on the fluid must be equal but oppositely orientated to that on the water droplet, this formula can then be used to determine the force on the column. Using the formulas:

$$P = m / A \quad \text{and} \quad m = \frac{F}{a}$$

we can calculate the resultant pressure of this force on the column, stopping it from rising unrealistically.

4 Implementation

The previous section presented the underlying physics of our modelling approach. This section explains several implementation details.

4.1 Particle System

Particles have a position, velocity, and mass but no volume. They can be affected by forces, either from outside the system or from other particles and can represent anything from rigid structures to fluids (as described earlier).

Our particle system is based on the work by Witkin [15] and was written as general as possible to allow for later extensions. All forces and environmental constraints are represented by interfaces that can be implemented as required. The system itself implements the interface required by the solver and as such represents itself as a point moving through $6n$ dimensional space where n is the number of particles and the position and velocity of each particles is represented by a 6-dimensional vector. The derivative of each particle's state vector $[x_1, x_2, x_3, v_1, v_2, v_3]$ is $[v_1, v_2, v_3, f_1/m, f_2/m, f_3/m]$ where f is the force acting on the particle. By representing the system as a point in $6n$ -dim space, all forces between particles can be applied simultaneously so that the system stays consistent for each time step.

Solving for a particular time step begins with the system calculating the initial forces acting on the particles. The particle system can then pass itself to the solver method of choice. Because the solver views $f(x,t)$ as a black box (i.e. doesn't know how the function is evaluated) it calls back for solutions at intermediately positions. At these points the system recalculates the forces and responds appropriately.

After the new positions and velocities are determined collision detection with the environment is performed. Panels defined in the environment file are used as boundary conditions to stop particles escaping. While any parallelogram is allowed as a panel it is important to note that the collision detection used here is only valid for rectangles. To find if a particle has passed a panel our algorithm checks first using the following formula which side of the panel the particle is on:

$$(P - V_1) \bullet n \leq 0$$

where P is the particles position, V_i is one of the vertices of the panel and the unit normal of the plane on which the panel rests. A further check is necessary to see if the particle also lies within the rectangle described by the panel. If the four conditions below hold true then the point is within the parallelogram.

$$(P - V_2) \bullet (V_2 - V_1) < 0$$

$$(P - V_2) \bullet (V_2 - V_3) < 0$$

$$(P - V_4) \bullet (V_4 - V_1) < 0$$

$$(P - V_4) \bullet (V_4 - V_3) < 0$$

Where P is the particle's position and V_1 through V_4 are the four vertices of the rectangle in anticlockwise order. This is a simplification of the equation to check which side of a line a point is. As the sign is the only thing we are interested in it is more efficient to disregard excess calculation.

A final check is made to ensure the particle is not 'legally' on the other side of the panel. By tracing the particles path back one time step and checking that the old position was 'correct' it can be assumed that the particle has indeed hit the panel.

As described in the previous section particles also have to be tested as to whether they have hit the main volume of water. Because columns are stored in a 2D array in their physical order it is very easy to determine which column a particle is directly above. The system's minimum x and z points are stored as global variables so it is simply a matter of finding the difference between the particle's position on the x/z plane and the system minimums and scaling by the resolution, i.e.

$$index_i = (P_i - min_i) \times resolution$$

where $index$ is the array index, P the particle's position, min the minimum value for the system and $resolution$ the number of columns per unit length. If the particle is outside the bounds of the column array or its vertical position (y coordinate value) is below all columns in the system the particle is destroyed. Otherwise it is absorbed into the volume model as described previously.

Finally each particle is checked to ensure it hasn't exceeded its lifetime after which the particle system clock is increased by the time step given and the system is again deemed stable.

Initially, when the particle system was being used exclusively, it was considered appropriate to try to model cohesion between the particles to mimic molecular interaction. The first attempt was to apply a simple spring with a limited area of effect. This spring forced particles towards an equilibrium situation. This then evolved through the application of the work done by Miller & Pearce [5] to that shown in figure 2. These forces were not found to be appropriate however as they made the fluid appear too viscous to be believable for water. They were also very slow.

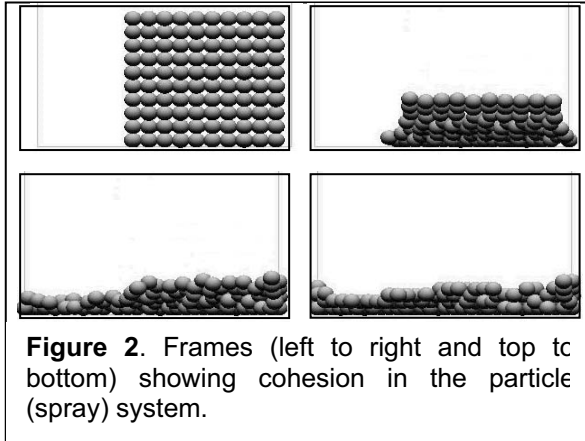


Figure 2. Frames (left to right and top to bottom) showing cohesion in the particle (spray) system.

4.2 Column System

As discussed above the volume model follows that used by Mould & Yang [8] closely. Before the system can be used columns must be generated from the environment specified and pipes must be created between all adjacent.

Using the same environment file format as that originally used for the particle system, each shape is cycled through checking if columns should be created. Spheres, used to represent rocks and obstacles are considered first. For each sphere in the environment a cross-section along the xz -plane is used. Taking the lower left corner of a bounding rectangle each possible column position is checked to determine if it is within the cross-section. For this an approximation is used where each corner is checked to see if it is inside, if two or more corners are inside then a column is created in that position. The base height of each column created is found by determining the height of the centre as it would be projected onto the sphere.

Each panel is also cycled through to create columns although, unlike spheres, each panel can serve one of three different purposes. If the area, on the xz -plane, is zero and the panel is not a source panel it is treated as a boundary panel for particle collision detection and is stored for use as such. Source panels also have no area but columns are still created along the line described by the first and second vertices. These columns have an initial height equal to the difference between the first and third vertices' y values which is kept regardless of outflows. However, most panels act in much the same way as all spheres by providing the ground environment for the column system.

Once again finding the bounding rectangle into which the projection of the panel onto the xz -plane fits, each possible column position is checked as to if a column should be created there. To do this opposite corners of the column position are tested as to whether they are left of the lines formed by V_1 and V_2 , V_2 and V_3 , V_3 and V_4 , and V_4 and V_1 . Each time a score is incremented for each corner to the left of each line, if both corners are outside a line then no column is created otherwise if the 'score' is above six (there was a maximum of two lines any corner was outside of)

then a column is created as appropriate. This method is only an approximation but serves well enough.

After each shape has been cycled through we now have an unordered array of columns. By sorting them into a 2D array so they are placed according to their physical position a much better representation of the environment can be made, and more efficient algorithms can be used later. If there are two columns for any position then either the higher or source column is taken.

Because the columns are stored in order the creation of pipes is rudimentary. Each column is cycled through creating pipes in the directions shown in figure 1 (b). As described in the model section all possible pipes are created and then checked for validity later instead of dynamically creating pipes in each time step.

To increment the system each pipe is cycled through with the intention of calculating flow. If a pipe's cross-section is less than 0 (i.e. the two cells don't overlap) nothing occurs, for those that do the flow and volume to be moved are calculated and the cells are updated accordingly. The user can choose the computational time step whereas the display time step is set to 0.005.

5 Results

Using hydrostatics and drawing on information from fluid dynamics the model described in this paper is capable of creating realistic effects for the simulation of rivers and waterfalls. The speed of the model depends largely on the number of columns needed to cover the area and the number of particles in the scene. For example the simple river scene shown in


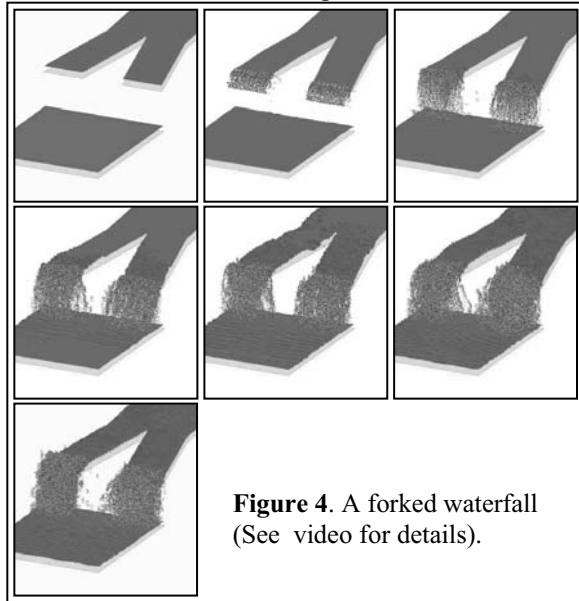


Figure 3. A Simple river with rapids around a large rock. See video file for more details.

figure 3 runs on a 1.4GHz PC without optimizations at 1 frame every 15 seconds with 13000 columns,

each with 4 cells and just over 3700 particles. Conversely the forked waterfall in figure 4 takes one and a half minutes per frame and has 27500 columns, each with 4 cells and 110000 particles.



The main drawback of the current model is the lack of viscous shear forces which is needed for creating eddies and visual effects such as shockwaves around rocks and near the edges of the river. Also we haven't yet incorporated spray from water particles hitting columns too hard. In reality this often causes a semi permanent area of white water and mist around the base of a waterfall.

6 Conclusion

Using hydrostatics and fluid dynamics our model is capable of creating recognizable effects for the simulation of rivers and waterfalls. The speed of the model depends largely on the number of columns needed to cover the area and the number of particles in the scene.

While the representation used here provides only an indication of the achievable effects and the current results are not yet rendered photo-realistically, the way in which the data from each frame is stored means that integration with a third party renderer would not be overly difficult. We are currently implementing more advanced implicit ODE solvers and current research indicates that this might lead to an order of magnitude improvement in speed which gets close to achieving interactive speeds.

7 References

- [1] Schachter, B. J., "Long crested wave models", *Computer Graphics and Image Processing*, 12:187-201 (1980).
- [2] Neyret, F. & Praizelin F., "Phenomenological Simulation of Brooks", *Computer Animation and Simulation*, pp 53-64, September (2001).
- [3] Reeves, W.T., "Particle Systems – A Technique for Modelling a Class of Fuzzy Objects". *Computer Graphics*, 20(3), pp 359-376, July (1993).
- [4] Sims, K., "Particle Animation and Rendering Using Data Parallel Computation", *Computer Graphics (Proceedings of SIGGRAPH '90)*, 24(4), pp 405-413, August (1990).
- [5] Miller, G. & Pearce, A., "Globular Dynamics: A Connected Particle System for Animating Viscous Fluids", *Computer & Graphics*, 13(3), pp 305-309 (1989).
- [6] Kass, M. & Miller, G., "Rapid, Stable Fluid Dynamics for Computer Graphics". *Computer Graphics (Proceedings of SIGGRAPH '90)*, 24(4), pp 49-57, August (1990).
- [7] O'Brien, J.F. & Hodgins J.K., "Dynamic Simulation of splashing fluids", *Proceedings of Computer Animation '95*, pp 198-205 (1995).
- [8] Mould, D. & Yang, Y.H., "Modeling Water for Computer Graphics". *Computer & Graphics*, 21(6), pp 801-814 (1997).
- [9] Chen, J.X. & Lobo, N., "Towards Interactive-rate Simulation of Fluids with Moving Obstacles by Navier-Stokes Equations". *CVGIP: Graphical Models and Image Processing*, 54(3), (1995).
- [10] Enright, D., Marschner, S. & Fedkiw, R., "Animation and Rendering of Complex Water Surfaces". *Computer Graphics (Proceedings of SIGGRAPH '02)*, pp 736-744 (2002).
- [11] Foster, N. & Fedkiw, R., "Practical Animation of Liquids". *Computer Graphics (Proceedings of SIGGRAPH '01)*, pp 23-30 (2001).
- [12] Thorton, E. & Guza, R., "Energy Saturation and Phase Speeds Measured on a Natural Beach". *Journal of Geophysical Research*, 87(12), pp 9499-9508 (1982).
- [13] Badger, W. & Banchero, J., *Introduction to Chemical Engineering*, McGraw-Hill (1955).
- [14] Stein, C. & Max, N., "A Particle-Based Model for Water Simulation", *Technical Report UCRL-JC-129378*, Lawrence Livermore National Laboratory, <http://www.llnl.gov/tid/lof/documents/pdf/233792.pdf>, visited on 16/08/2003.
- [15] Witkin, A., "Particle System Dynamics". *SIGGRAPH '94 Course Notes* (1994).