

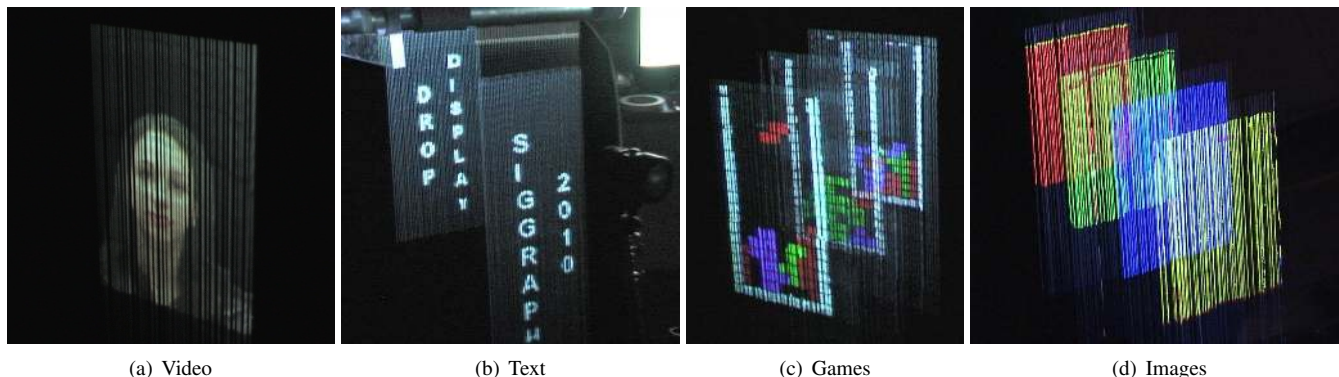
# A Multi-Layered Display with Water Drops

Peter C. Barnum

Srinivasa G. Narasimhan

Takeo Kanade \*

Robotics Institute, Carnegie Mellon University



**Figure 1:** By using precisely controlled valves and a projector-camera system, we create a vibrant, multi-layered water drop display. The display can show static or dynamically generated images on each layer, such as text, videos, or even interactive games.

## Abstract

We present a multi-layered display that uses water drops as voxels. Water drops refract most incident light, making them excellent wide-angle lenses. Each 2D layer of our display can exhibit arbitrary visual content, creating a layered-depth (2.5D) display. Our system consists of a single projector-camera system and a set of linear drop generator manifolds that are tightly synchronized and controlled using a computer. Following the principles of fluid mechanics, we are able to accurately generate and control drops so that, at any time instant, no two drops occupy the same projector pixel's line-of-sight. This drop control is combined with an algorithm for space-time division of projector light rays. Our prototype system has up to four layers, with each layer consisting of an row of 50 drops that can be generated at up to 60 Hz. The effective resolution of the display is  $50 \times \text{projector vertical-resolution} \times \text{number of layers}$ . We show how this water drop display can be used for text, videos, and interactive games.

**CR Categories:** I.3.1 [Computer Graphics]: Hardware architecture—Three-dimensional displays; B.4.2 [Input/Output and data communications]: Input/Output devices—Image display

## 1 Introduction

Water drops are transparent and serve as tiny fish-eye lenses. These lenses collect light and refract it out across a wide range of angles and can be among the brightest elements in an environment. By taking advantage of water drops' tendency to refract light and by illuminating collections of drops, vibrant displays can be created. Water drop displays have been used to show anything from abstract patterns to movie-theater quality video. These displays have either

been high-resolution 2D [Araki et al. 1991; Palovuori and Rakkolainen 2004], or low resolution 3D [Eitoku et al. 2009]. In this work, we have created a high-resolution, multi-layered (2.5D) drop display that can show images, videos, or games.

To create the display, we synchronized a projector, a camera, and multiple linear drop-emitter manifolds. The first step is to accurately generate layers of drops. Each layer is created by a linear manifold of emitters, controlled by a single solenoid valve. When the valve is actuated, the manifold emits a row of water drops that freely fall down a vertical plane. The timings of the drop emissions from different manifolds are offset, so that at any instant in time, no two drops occlude each other from the point of view of a projector.

The second step is to find the drops' positions and illuminate them with the correct colors. The camera captures several images of the locations of the falling drops. Based on the detected drop locations, the projector quickly switches images, so that drops from each layer are illuminated with the correct colors. Figure 1 shows example usages. The display can show static images like the text in Figure 1 (b) or the solid-color images in Figure 1 (d). It can also play pre-recorded videos, like the video of the woman in Figure 1 (a), or even interactive content, as with the Tetris variant in Figure 1 (c).

Previous 2D and sparse 3D displays offer an enjoyable experience, but they do not offer functionality beyond a touchscreen or a laser-light show. But a high-resolution 3D drop display would be uniquely useful. Since water is not solid, both the outside and the inside of the display volume could be directly touched, opening up new interaction possibilities. By stacking more layers, our 2.5D display could scale to become a high-resolution 3D display.

## 2 Previous Work

Water drop displays can be compared by asking two questions. First, how accurately can the drops' locations be predicted or controlled? A flat 2D display requires no control, but a 3D display with drops occluding each other at different times needs precisely positioned drops. Second, how quickly and accurately can the light source be controlled? The faster and more accurate the light source, the more tolerance the display has to drops occluding each other.

The simplest drop displays control only the light source. Fog screen displays [Araki et al. 1991; Palovuori and Rakkolainen 2004; Lee et al. 2007] use a sheet of microscopic droplets and a standard projector to create high-resolution, walk-through projector screens.

\*e-mail: {pbarnum,srinivas,tk}@cs.cmu.edu

But because they have no control over individual drops, they require a different projector for each screen.

In addition to light-control, some displays selectively turn on and off fountains of water, such as Disney’s Fantasmic show [Disney 1992-2010] and the Submerging Technologies [Dietz et al. 2006] pieces. Such limited water control allows for interesting interactions and large-scale performances. The purpose of these displays is to create specific artistic effects, so they should not be evaluated only with quantitative metrics. However, this limited water control does not allow complexity much beyond that of fog screens.

Instead of controlling the light source, some displays, such as [Pevnick 1981], use a constant light source and arrays of individually controlled drop emitters to print out falling images. This has the advantage of spatial scalability. But as the drops are constantly falling, individual images can only be shown temporarily, and this display cannot be used for video.

As shown by the previous examples, fun and colorful water displays can be made by controlling only the water or only the illumination. But the most versatile types use a combination. Some such displays use uncontrolled drops, but detect the drops’ positions in real time. Section 3 of [Barnum et al. 2009] shows a method of detecting, tracking, and illuminating individual water drops. But due to the slow reaction time of the camera-projector system, it is limited to only a few drops at a time, and thus cannot create coherent images. But as suggested in the patent [Perlin and Han 2006], a high-speed adaptively-controlled laser could show arbitrary 3D shapes in particulate media (such as dust). It will be interesting to see if and how this can be implemented.

Instead of finding drops in real time, other displays use precise drop emission control and a synchronized projector. Eitoku *et al.* [Eitoku et al. 2006; Eitoku et al. 2009] create a low-resolution 3D display by emitting a 2D array of water drops that are illuminated by an LCD projector. This display can be used to create colorful abstract patterns and shapes. But due to limitations in drop and illumination control, it only has a resolution of about  $16 \times 16 \times 8$  at 7.5Hz, and requires manual computation of the drop positions.

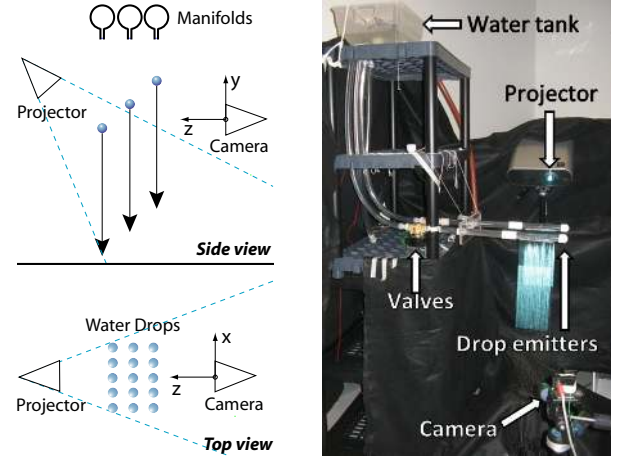
### 3 Overview of the Display System

Figure 2 shows our setup in action. The left subfigures illustrate the positions of the projector, the camera, and the drop-generation manifolds. Each manifold contains fifty drop emitters arranged linearly, and is controlled by a solenoid valve. The valve cycles at up to 60 Hz, releasing up to 3000 drops every second for each layer. The drops fall fast and are quickly replaced, and to the naked eye, they appear as a row of streaks along a vertical plane. But because each manifold releases drops at slightly different times, no drop ever occludes another from the point of view of the projector.

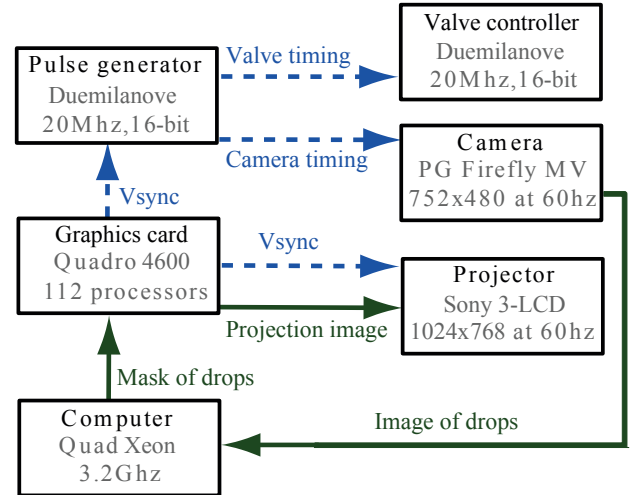
Consider the order that the drops enter the projector’s field of view (the top dotted line in the side view of Figure 2). The leftmost drop has already crossed it, next the middle one will, followed by the rightmost. The projector starts displaying an image when the first drop crosses, then switches images as the second drop crosses, and so on. For example, if the drops are offset by  $1/10$  seconds and we do not consider perspective effects, then the projector has to switch frames at 10 Hz. The optimal timing offsets has drops spaced as far apart from each other as possible. For  $N$  manifolds operating at  $S$  Hz, the drops should be separated by  $1/(SN)$  seconds. A simple bound on the necessary projector refresh  $P$  is therefore:

$$P \geq SN \quad (1)$$

The equation predicts that a single 60 Hz projector could illuminate 5 layers at 12 drops per second, 2 layers at 30 drops per second, or trivially, one layer at 60 drops per second. This equation is a



**Figure 2:** Diagrams and a photo of the drop display setup. By emitting drops and illuminating them with a Sony VPL-HS51A 3LCD projector, we can create multi-layer water drop displays. (The left subfigures show the setup from the side and from the top. The photo is taken from slightly behind the Point Grey Firefly MV camera.) The water head is about 0.8m above the valves, and is regulated by a Rule 360 Bilge Pump and a float switch in the water tank.



**Figure 3:** A chart of the data and synchronization flow. The dotted lines represent the synchronization signal, which starts from the graphics card, and travels to the projector, camera, and valve controller. The solid lines represent data flow. During the calibration step, the camera detects the locations of the drops and sends the images to the computer. The computer detects and stores the drop locations. After calibration is complete, the graphics card rapidly sends images to the projector, which illuminates the drops. (The pulse generator and the valve controller are built around separate Arduino Duemilanove microcontrollers.)

reasonable first approximation, although we have included a tighter upper bound in Appendix A, that accounts for uncertainties in drop emission times and finite projector switching time.

Determining the locations of the drops at any time instant and projecting to hit them requires precise design, synchronization and control. The flowchart in Figure 3 shows the passage of timing signals and data (images) between different components of our system. To start the process, the graphics card sends a V-sync timing pulse to the projector and to a custom pulse generator. At the same

time as the projector begins displaying an image, the pulse generator can signal the valve controller to begin releasing the drops from the first manifold, and can signal the camera to begin acquiring an image. Because the projector, camera and valves are all synchronized, there are only a small set of drop-locations for all projector refresh cycles. From Equation 1, we see that for a projector running at  $P$  Hz and drops falling at  $S$  Hz, there are only  $P/S$  possible states. This means that capturing  $P/S$  images from the camera and warping these images to the projector reference frame tells us which projector pixels correspond to which drops at which times. And because these locations are precisely known, it is straightforward to direct the projector to illuminate them. The individual components for creating this precise setup are detailed in the following sections.

## 4 Accurate Generation of Drops

Accurately generating well-defined and stable drops at precise time instants is crucial to our display. For this, we make the following design choices based on basic principles of fluid mechanics.

### 4.1 Drop sizes

Coherent water drops can exist stably with diameters between about  $1\mu\text{m}$  and  $5\text{mm}$  [Pruppacher and Klett 1997]. Different sized drops cause different visual effects. The smallest, microscopic drops cause fog and mist. Larger drops are individually visible, and in nature, are raindrops. Both microscopic drops and macroscopic drops have successfully been used to create displays. Microscopic drops (mist or fog) have been used to create high-resolution flat displays with a projector [Araki et al. 1991]. Using multiple planes of mist could allow for a semblance of 3D but would result in a loss of contrast due to light scattering as the light passes through each plane.


Although a display with macroscopic water drops will have fewer voxels, the ability to control drops in space and time makes it the best choice for our system. By using millimeter-sized drops, we trade 2D resolution for depth and create a multi-layered display. Such drops are less effected by air currents and are hence easy to handle. Even so, we must be careful about choosing their exact size — if they are too large, they will break apart or cause satellite drops. If they are too small, they can adhere to the nozzle and coalesce with each other, leading to irregular timing. As discussed in [Ambravaneswaran et al. 2002], different combinations of size, ejection velocity, and fluid characteristics effect the ability to create usable drops. From their work and our experimentation, we find that drops approximately 2mm in diameter provide a good combination of ease-of-handling and structural stability.

### 4.2 Fluid mechanics of Drop Generation


One option for our display is to control every emitter independently using a dedicated valve [Pevnick 1981]. Though providing accurate drop generation, it would be costly for large numbers of emitters. Instead, we choose to control a linear manifold of emitters with one valve (Figures 4 and 5). The disadvantage is that the mechanics of controlling the formation of a row of drops with one actuator is more complex, and the drops may behave less predictably.

**The emitters:** First, we must design the individual drop emitters. As discussed in [Ambravaneswaran et al. 2002], important factors in the accurate emission of drops are the *Ohnesorge number* (which relates the viscosity to inertia and surface tension), the *Weber number* (which relates inertia with surface tension), and the thickness and hydrophobicity of the emitter walls.

Since we are using pure water, the main factors we can vary are the materials, and the wall thickness and diameters of the emitters. We choose Type 304 stainless steel, as it is fairly hydrophobic and its strength allows us to minimize the wall thickness. Hydrophobicity is increased by applying a Rain-X coating to the emitters' exteriors. Based on the results of [Ambravaneswaran et al. 2002] and our



**Figure 4:** An illustration and photo of a water drop emitter manifold. When charged, the solenoid pulls up the valve, allowing water to pass from the high-pressure source to the feed. This will pressurize the feed, causing drops to be formed at each emitter. The valve is a generic 3/4 inch, 12VDC, 2 Amps, normally-closed diaphragm solenoid valve. Its Cv flow coefficient is about 5, although we only turn on the solenoid for about 10ms, so the valve never completely opens and the achieved flow rate is much less.



**Figure 5:** These photos show two 4ms exposures, taken 1/30th of a second apart. The drops were emitted from a linear manifold, as shown in Figure 4. The valve was actuated at 30hz, so the rows are spaced by 1/30th of a second. Although the drops are not in completely straight horizontal lines, each line is nearly identical. If a manifold is emptied and refilled with water, it will have a different line pattern. But as demonstrated in these images, there is almost no change over the short term.

own experimentation, creating emitters with an inner diameter of 1.27mm allows predictable and fast drop generation, with few or no satellite droplets (Figure 5 and our video).

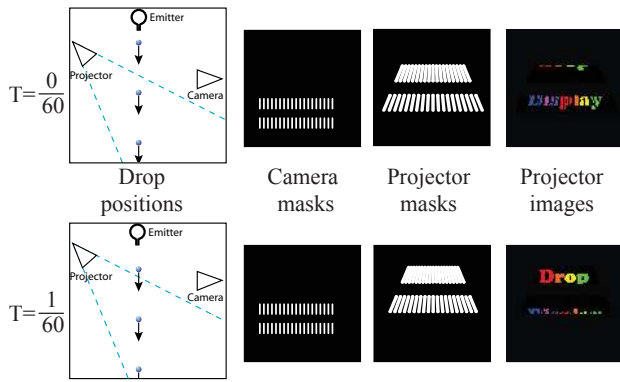
**The feed:** For drops to be released in the same way by all emitters, the feed needs to have similar pressure at all emitters. The fluid flow in a closed cylindrical pipe can be predicted using well-known principles of fluid mechanics [Munson et al. 2009]. Accordingly, assuming laminar flow and no vertical pressure variation, the water pressure at each emitter of a pipe manifold depends on the distance between the emitters  $l$ , the flow rate  $Q$ , and the diameter of the pipe  $D$ . The pressure difference between two emitters is given by:

$$\frac{128\mu l Q_e}{\pi D^4} \quad (2)$$

where,  $Q_e$  is the fluid flow out of the emitter further from the source, and  $\mu$  is the dynamic viscosity of water. To minimize the pressure difference between emitters, we can either reduce the distance  $l$  between them, or increase the diameter  $D$  of the pipe. Therefore, to create our emitter manifolds, we use a large diameter tube (19mm) and place the emitters as close as possible (2.5 mm).

**Electronics for controlling the valves:** Creating 10 drops per sec-





**Figure 6:** A summary of the drop display process, using drops being emitted at 30hz and a projector refreshing at 60hz. The water drops fall a short distance in each 1/60th of a second. The camera captures images as they fall, which are thresholded to create binary masks. The camera mask creation is performed before any content is displayed. During runtime, the camera masks are warped with a homography to the projector coordinate frame. The warped masks are then used to create the final projection image.

ond from 50 emitters requires a flow rate of  $10 \times 50 \times [\text{volume of a drop}] \approx 1.7 \times 10^{-5} \frac{m^3}{s}$ . Even slight variations in the flow rate can prevent correct drop emission. This means that the opening and closing of the solenoid valves must be accurate to within a few microseconds. Using a dedicated microcontroller, each solenoid is controlled by a MOSFET. The microcontroller is synchronized by a timing signal from the pulse generator (Figure 3).

## 5 Space-Time Division of Illumination

The ability to easily design visual content is key to a display’s usefulness. We have described how to densely fill planes with non-occluding water drops. Similar to [Barnum et al. 2009], we will allow the user to “paint” the drops from the point of view of the camera. The user will select images or movies to display on each layer, and the system will automatically determine how to control the projector to display them, with no further user input. This involves three steps. First, the correspondence between camera rays and projector rays is computed. Second, the camera is used to determine where the drops of each layer are at every time instant. Finally, using the detected image locations, the projector is controlled in real-time to display images on all layers.

**Projector-camera calibration:** To compute the correspondence between projector and camera rays, we first display and capture a checkerboard pattern at each layer. The user then selects a few point correspondences between the captured image and the original pattern, and a homography is computed for each layer, which maps image coordinates to projector coordinates.

**Drop location estimation:** As the valves, projector and camera are synchronized (Figure 3), the display is created as illustrated in Figure 6. For example, let us consider a single manifold emitting drops at 30 Hz, illuminated by a projector operating at 60 Hz. The state of the drops will repeat every 1/30 seconds, during which time two images can be projected. Recall that for a projector operating at  $P$  Hz and drops falling at  $S$  Hz, there are  $P/S$  unique projector states. To determine where the drops are in each projector state, a camera acquires images at 60 Hz, shown in the second column of Figure 6. These binary images contain the locations of the drops during the time the projector can display its images.

**Real-time synthesis of display content:** Once the projector-

camera calibration and drop location estimation are done, the display is ready to run. Using an NVIDIA Quadro 4600 GPU and the homographies computed in Section 5, the binary drop masks are transformed to the projector coordinate frame (third column of Figure 6). The user selected image (Drop Display, in this example) is also warped to the projector coordinate frame, and is masked. When the projector displays the warped and masked image, the drops will be illuminated, making the text appear. Because the projection images are computed and displayed in real-time, the user can modify them as the display operates, and any of the layers of drops could even serve as an interactive touch screen.

## 6 The Display in Action

Figure 1, Figure 7, and our video show several example displays. We have demonstrated displays with up to four layers. Each layer is made of rows of 50 drops, and the projector resolution is 1024x768.

The drops are spatially sparse, but the layers appear solid to the human eye, with some flicker. As a camera integrates temporally differently than an eye, we have been unable to reproduce the visual effect in our videos, but no one at our demos has found the flickering disturbing. Although it is possible to run individual layers at up to 60 drops per second, we found that the decrease in perceived flicker is marginal above 15 drops per second. The effect is most prominent during eye or camera motion. When the camera moves (as shown in our video), different pieces of a layer appear to be in slightly different positions, causing breaks in the images. Interestingly, due to the skill of the human visual system, these position shifts are not seen in real life.

The masking and warping takes approximately 5ms, leaving about 10ms free to load video frames and render scenes. The pre-masked-and-warped image for each layer is stored as a DirectX Surface, so any application that can render to a surface can be seamlessly used.

The first example in Figure 7 (a) shows one layer at 60 drops per second, with a video of a woman from the VidTIMIT dataset [Sanderson 2008]. Figure 7 (b) has two layers with text. The left-most image demonstrates that the text is still bright enough to be seen with the overhead lights turned on.

Although the display is brightest when the viewer faces the projector, Figure 7 (c) demonstrates that the display has a wide viewing angle. It also demonstrates the display’s accuracy. Every projector pixel has to switch colors at exactly the right time, making this example particularly difficult. Any errors in projector-camera calibration or drop emission timing will be instantly visible. Figure 1 (d) shows a similar four-layer example at 5 drops per second.

Figure 7 (d) and (e) demonstrate how we can generate content on-the-fly. Figure 7 (d) is an aquarium simulator, with swimming fish and waving plants. Figure 7 (e) is a version of Tetris played on multiple layers. The player uses a remote presenter mouse to move and rotate pieces, and to move pieces between layers.


## 7 Conclusion

By using a camera, a projector, and several water drop emitter manifolds, we have created a 2.5D water drop display that can be used to show pictures and video on-the-fly. The drops are generated in sync with the camera and projector, therefore the camera can be used to automatically discover the drop locations. This allows users to easily create content without requiring extensive manual calibration.


The main limitation of this work is that it only has a few layers. As discussed in Section 3 and Appendix A, a faster projector could allow for a display with many more layers. But to create a true 3D drop display would also require more precise drop control. The




(a) A single layer with a video of a woman from the VidTIMIT dataset (60 drops per second).




(b) Zooming in on text (15 drops per second): Because the drops oscillate as they fall, they appear as wiggly lines when viewed close up. But because they are generated with the same initial oscillation, each drop distorts in the same way as it falls, and the wiggles appear stationary.



(c) Red-Green-Blue (10 drops per second): The display is brightest when viewed from directly opposite the projector (leftmost image). But the colors are still clearly visible at an extreme angle. These images were taken with fixed camera settings, and the intensity change is much less evident to the adaptive human eye.



(d) Aquarium simulator (10 drops per second): The projector images are created in real time and we can generate content on-the-fly. This aquarium simulator has swimming fish and waving plants. To create the leftmost image, we took the pixelwise max for the right three images (30FPS, 1/30 second exposure). The human eye temporally integrates the falling drops, so the display appears like the left image. But the other images demonstrate that the drops are actually spatially sparse.



(e) 2.5D Tetris (10 drops per second): This twist on the classic game demonstrates interaction with the drop display. Only one piece is dropped at a time, but the player can choose which layer to use it in. The rightmost image shows a player using a remote presenter mouse to move and rotate the pieces.

**Figure 7:** Several examples of different images and movies displayed on layers of water drops.

simplest way to increase accuracy would be to better flush air bubbles from the valves and manifolds. (The bubbles are caused by air dissolved in the water and by leaking at the emitters' water/air interfaces). Alternately, drops could be created on demand via other techniques [Le 1998], although many drop-on-demand applications may only be usable for micrometer-scale droplets. Finally, drop streams could be created by disturbing a jet, as summarized in [Frohn and Roth 2000]. Using these alternate techniques could greatly increase the number of drops in the environment, allowing for dense 3D, or perhaps even immersive water drop displays.

## A Tighter upper bound on scalability

By Equation 1, the required projector refresh rate scales linearly with the number of layers and the number of drops emitted per second. But this assumes that the drops are emitted with identical initial velocities, and the projector can switch frames instantaneously. A more accurate upper bound considers drops with different initial velocities and finite frame switching time.

Let us assume that the projector is orthographic, and the drops are points. For an initial velocity  $v$ , the vertical position  $y$  of a drop at time  $t$  is given by the free-fall equation:

$$y(t, v) = vt - 4.9t^2 \quad (3)$$

We need to determine if the beginning location  $B$  of a drop on one layer intersects with the end location  $E$  of a drop on the next layer. Let the projector refresh rate be denoted by  $P$ , the current projector refresh cycle by  $\tau \in \mathbb{N}$ , the LCD switching time by  $\rho$ , the initial velocity of drops from the  $i$ -th emitter on the  $j$ -th layer by  $v_{i,j}$ , and the time offset for the  $j$ -th layer by  $o_j$  (for example, if  $o_1 = 0$  and  $o_2 = 0.1$ , then the second manifold releases drops 0.1 seconds after the first). Typical values for these quantities are  $P = 60\text{hz}$ ,  $\rho = 0.01$  for an LCD projector, and  $v_{i,j} \approx 0.4\text{m/s} \pm 0.05\text{m/s}$ . Therefore, for the  $k$ -th drop emission, the beginning location  $B$  of a drop on one layer and the end location  $E$  of a drop on the next layer can be computed by substituting into the free-fall equation  $y(t, v)$ :

$$B = y\left(\frac{\tau}{P} - o_j - \frac{k}{s}, v_{i,j}\right) \quad (4)$$

$$E = y\left(\frac{\tau + 1}{P} - o_{j+1} - \frac{k}{s} + \rho, v_{i,j+1}\right) \quad (5)$$

To check if any drop on any layer intersects the same projector pixel at any time, for all  $j$  such that  $o_j < o_{j+1}$ , we minimize the following expression:

$$\operatorname{argmin}_{o_1, \dots, o_N} \left\{ \sum_{j=1}^{N-1} \left[ S\left(\frac{\tau}{P} - o_{j+1}\right) \sum_{k=0} \max(0, B - E) \right] \right\} \quad (6)$$

If the result is greater than zero, then there are drops that cannot be illuminated different colors.

An LCD projector switches at  $\rho = 0.01$  and has a projector refresh rate of  $P = 60$ . Therefore, according to the model, the display scales to 3 layers at 10 drops per second, or 2 layers at 15 drops per second. We have successfully implemented 3 layers at 10 drops per second and 2 layers at 15 drops per second, which validates this theoretical estimate.

If instead of the LCD, we used one 60hz DLP for each color channel, then the frame switching time  $\rho$  would be on the order of microseconds. With three  $P = 60\text{hz}$  DLPs, we could have 4 layers at 10 drops per second, or 3 layers at 15 drops per second.

There are also projectors, such as the Pico 2, that are user controllable at up to 1440hz. These high-speed projectors could allow for displays with additional layers. For  $P = 1440\text{hz}$  switching projectors, we could have 17 layers at 10 drops per second, or 11 layers at 15 drops per second. In this case, the main bottleneck is the variance in velocities and not the projector refresh rate.

## Acknowledgments

This work was supported in parts by an Okawa Foundation Research Grant, ONR Grants N00014-08-1-0330 and DURIP N00014-06-1-0762, and NSF CAREER IIS-0643628. We would also like to thank Lee Weiss, Shelley Anna, and Lynn Walker for help with the fluid mechanics, Prasanna Velagapudi for help with the electronics, and Larry Hayhurst for help with precision machining techniques. We would also like to thank Scott Satkin for ideas on interesting things to display.

## References

- AMBRAVANESWARAN, B., WILKES, E. D., AND BASARANA, O. A. 2002. Drop formation from a capillary tube: Comparison of one-dimensional and two-dimensional analyses and occurrence of satellite drops. *Physics of Fluids* 14, 8, 2606–2621.
- ARAKI, T., KAWAMATA, F., OGINO, M., MIYAGAWA, H., KAWAMATA, T., WATANABE, M., AND MIYASHITA, K. 1991. US Patent 5,067,653: Screen forming apparatus and method.
- BARNUM, P. C., NARASIMHAN, S. G., AND KANADE, T. 2009. A projector-camera system for creating a display with water drops. In *IEEE International Workshop on Projector-Camera Systems*.
- DIETZ, P., HAN, J. Y., BARNWELL, J., WESTHUES, J., AND YERAZUNIS, W. 2006. Submerging technologies. In *SIGGRAPH Emerging Technologies*.
- DISNEY. 1992-2010. Fantasmic!
- EITOKU, S., TANIKAWA, T., AND SUZUKI, Y. 2006. Display composed of water drops for filling space with materialized virtual three-dimensional objects. In *IEEE conference on Virtual Reality*.
- EITOKU, S., NISHIMURA, K., TANIKAWA, T., AND HIROSE, M. 2009. Study on design of controllable particle display using water drops suitable for light environment. In *ACM Symposium on Virtual Reality Software and Technology*.
- FROHN, A., AND ROTH, N. 2000. *Dynamics of Droplets*. Springer.
- LE, H. P. 1998. Progress and trends in ink-jet printing technology. *Journal of Imaging Science Technology* 42, 49–62.
- LEE, C., DIVERDI, S., AND HÖLLERER, T. 2007. An immaterial depth-fused 3D display. In *ACM Symposium on Virtual Reality Software and Technology*.
- MUNSON, B. R., YOUNG, D. F., YOUNG, D. F., OKIISHI, T. H., AND HUEBSCH, W. W. 2009. *Fundamentals of Fluid Mechanics*, 6th ed. Wiley.
- PALOVUORI, K., AND RAKKOLAINEN, I. 2004. US Patent 6,819,487: Method and apparatus for forming a projection screen or a projection volume.
- PERLIN, K., AND HAN, J. Y. 2006. US Patent 6,997,558: Volumetric display with dust as the participating medium.
- PEVNICK, S. H. 1981. US Patent 4,294,406: Program controllable free falling water drop fountain.
- PRUPPACHER, H. R., AND KLETT, J. D. 1997. *Microphysics of Clouds and Precipitation*, second revised and enlarged ed. Kluwer Academic Publishers.
- SANDERSON, C. 2008. *Biometric Person Recognition: Face, Speech and Fusion*. VDM-Verlag.