DIGITAL MICROFLUIDICS DROPLET DISCOVERY AND IDENTIFICATION THROUGH USE OF HOUGH TRANSFORMS

By

William Fergus

A Thesis Submitted to the Graduate

Faculty of Rensselaer Polytechnic

Institute in Partial Fulfillment of the

Requirements for the Degree of

MASTERS OF SCIENCE

Major Subject: COMPUTER SCIENCE

Approved:

____

Robert Linhardt, Thesis Advisor

____

Chris Bystroff, Thesis Advisor

Rensselaer Polytechnic Institute
Troy, New York

May 2011

TABLE OF CONTENTS

LIST OF FIGURES

ACKNOWLEDGMENTS

The author wishes to express sincere appreciation to Drs. Linhardt and Bystroff for their assistance in the preparation of this manuscript and for their support during research. In addition, I would like to thank Jeff Martin, Julie Beaudet, Justin Gullotta, and John Schwartz for all of their help in the research and implementation stages.

GLOSSARY

Artificial Golgi. A digital microfluidics device developed at Rensselaer Polytechnic Institute.[1] All results were collected on this device.

Canny Edge Detector. One of the popular methods for finding edges in an image, discovered by John Canny.[2]

Capacitance. The ability of a body to hold an electrical charge. In digital microfluidics, capacitance can be used to determine the location of water droplets.[1]

Click Chemistry. A common experiment to be performed on a digital microfluidics device. In this style of chemistry, a compound is "clicked" onto a nanoparticle, resulting in a droplet of "product" and several droplets of different kinds of "reagent" with which the product is intended to react.[1]

Digital Microfluidics. A lab-on-a-chip system in which the property of electrowetting (EWOD) is utilized to combine droplets of solution to cause small-scale reactions.[1]

Electrowetting. The process of modifying the hydrophobic/hydrophilic properties of a surface to induce movement of aqueous solution.[1]

Hough Transform. A method for determining the presence or absence of a line or other geometric construct through voting arrays constructed over the image's pixels as a group.[3]

Ring Light. A light attached to the aperture of the camera, suspended above the surface to be observed. Ring lighting can cause a great deal of unwanted glare if not properly configured.

ABSTRACT

This thesis in computer science outlines the problems inherent in using classical computer vision algorithms on a digital microfluidics device while making few assumptions about the nature of the aqueous solution being utilized on the device. The difficulty arises from attempting to properly locate single or potentially multiple distinct droplets of unknown color, transparency, and size suspended in a single contiguous background pool of silicone oil. A multi-step algorithm is presented that is capable of recognizing such droplets with good success. Several simple changes can be made to the system to improve accuracy and speed while still staying within the constraints and assumptions of the original system. The algorithm is presented in three distinct steps: a basic pre-processing step, a search-space reduction step to improve speed and remove false positives, and a circle-detection step. This algorithm is shown to work reasonably well in a poorly constructed system and to produce much better results with certain additions and changes to that system.

1. INTRODUCTION

1.1 The Artificial Golgi Device

Microfluidics and lab-on-a-chip technologies have recently made a big splash in the fields of chemistry, biochemistry, and electrical engineering due to their ability to enable chemical reactions on the microscale. This capacity provides higher control of reactions and also allows them to be performed with fewer reagents, reducing their costs.[6] Additionally, microfluidics allows for great amounts of automation, reducing and possibly eliminating the need for an experimenter to be present throughout the chemical reaction.

1.1.1 System Automation

To automate these systems, as with most physical systems, two important subsystems need to be developed. First, there needs to be a controller, which commands or directs the system and is in charge of telling the system how to function. Programming the controller is usually straightforward. Second, there needs to be an observer, which checks to make sure that the commands sent by the controller were actually performed. This observer is not as easy to implement. In the case examined in this thesis, there are some integral restraints to the system that make the problem of programming the observer increasingly difficult to solve.

1.1.2 Droplets and Motion

Digital microfluidics, also known as the electrowetting-on-dialectric (EWOD) effect, typically involves a system set up as shown in Figure 1 (below). A droplet of solution is placed on a surface consisting of an electrode with a hydrophobic insulator coated on top of it. The hydrophobic property of the surface causes the solution droplet to bead up on the surface as shown on the left side of Figure 1, in order to get as far away from the surface as possible. If a ground is also present touching the droplet (usually provided by suspending it over the surface), when a voltage is applied to the electrode beneath the surface, the surface becomes more hydrophilic and the droplet changes its angle of contact, appearing to "stick" to the surface.

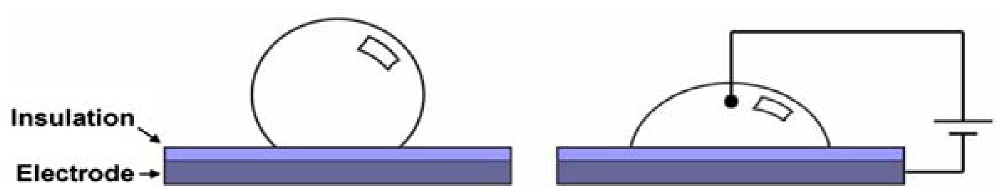

Figure 1: EWOD Effect, Adapted from Fair [5]

If the droplet is slightly larger than a single electrode and is placed above a pattern of electrodes, and if a voltage is applied to specific electrodes, the droplet will have different angles of contact based on the hydrophilic/hydrophobic properties of the electrodes in the pattern. This effect is shown in Figure 2, where a voltage has been applied to the right of the droplet but not to the left. Because of this effect, the droplet will shift towards the more hydrophilic electrode (on the right) and away from the hydrophobic electrode (on the left), causing it to move in the direction of the arrow until it is centered over the hydrophilic electrode.[1]

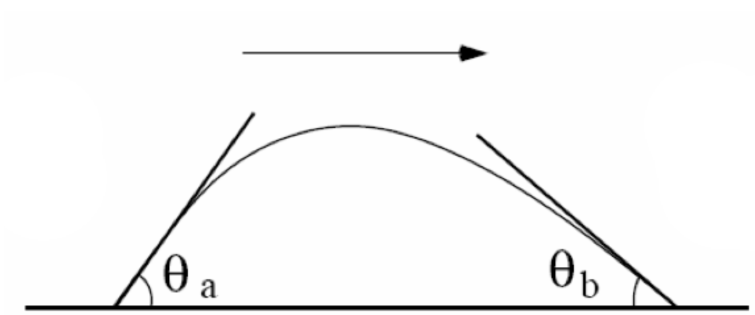

Figure 2: Droplet Movement, from Martin [1]

1.1.3 Programming the System

The controller for a lab-on-a-chip system is rather straightforward. All that is required is a program that takes as input an experiment plan provided by a human experimenter and the current location of all of the required reagents. This program then determines which reagents to mix when and where, and sends voltages to the appropriate electrodes to move the droplets as required. While the higher-level programming of the user-interface and input-parser to facilitate easy communication between the experimenter and the controller are clearly not trivial, they are straightforward to program once a convention has been chosen. They should in any case be tailored to the individual users. The lower-level programming, on the other hand, is simple, as the only important part is for the program to send a command to a complicated set of voltage-gates to turn them on and off at specific intervals.

1.1.4 The Observer's Task

As in most physical systems such as automated bank tellers, nuclear power plants, and license-plate scanners, it is important to make sure that the system is performing in the way the controlling program believes it should be. This functionality is typically implemented as a feedback loop: an observer takes in some sort of feedback signal, determines whether it is correct or not, and chooses an appropriate action to correct the system as needed. Such an observer is also necessary in the case of a lab-on-a-chip. The surface of the device gets slowly dirtied by residues from the reagents that pass over it, requiring more voltage to be applied to cause the same results. Because of this, sometimes when an electrode is activated, it does not have enough "pull" to center the droplet on it before it is deactivated and the next electrode activated. This can cause a droplet to be left behind, possibly making it miss a reaction or causing another droplet to collide with it, causing an unwanted reaction. The observer must recognize this problem and deal with it by interrupting the controller's program and resending the command that did not complete successfully, possibly with a higher voltage, until it is successful. This task can be broken down again into three parts: receiving the feedback signal, determining whether something is wrong, and sending the appropriate signal to fix it.

1.2 Feedback Data

The first task of the observer is to gather or receive the feedback signal. To do this, the observer must have some way of sensing the current state of the system accurately and without obstructing the general operation of the system until something goes wrong. One solution proposed in the literature is to measure the capacitance between the surface and the grounding slide at each electrode. The capacitance reading will be an order of magnitude higher if there is a droplet of solution between the surface and the grounding slide.[1] However, to do this, a capacitance reader needs to be added to the system and capacitance readings need to be performed over all electrodes at every step. Due to the nature of capacitance and the value of the equipment in question, the voltage must be turned off before the capacitance reader is connected to perform the readings. Clearly, this obstructs the system unnecessarily by halting the general operation of the system even in the case where nothing is wrong. Another solution is vision. By suspending a camera over the system, information can be gathered about its current state without interrupting the flow of commands from the controller to the system. This is the solution pursued in this thesis.

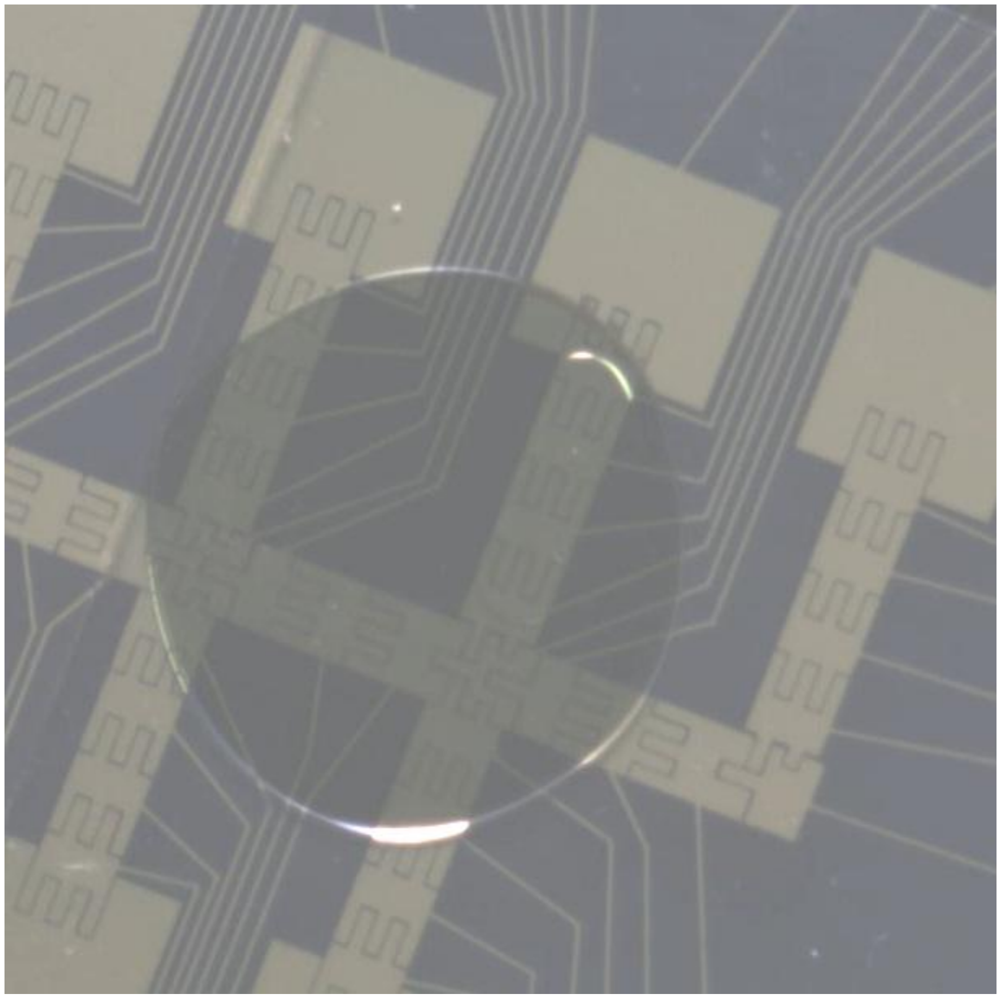

Figure 3: Sample Image Data

1.2.1 Problems with Vision Data

Unfortunately, several factors make it difficult to interpret the data obtained using the camera. Figure 3 shows an example of good image data that can be used to illustrate this point. In the background of the picture you see one of the many possible patterns of a digital microfluidics device. This particular device, called the Artificial Golgi, was developed at Rensselaer Polytechnic Institute. It prevents small droplets from ever being stranded over only one electrode by giving all neighboring electrodes "teeth" into each other. This design allows any droplet to be manipulated by the pulling force of at least two electrodes at all times.[1] The foreground of the image is dominated by a very large drop of an inert, transparent, colorless oil medium used by this device to facilitate the movement of droplets. Visible near the center of the image is the upper-right edge of the solution-droplet. To perform as intended, the observer must successfully find this droplet. The droplet's most prominent edge has been highlighted well in this picture by the lighting, a trait that makes this image a good example of the data. Occasionally, droplets will be more obvious due to the specific coloration or opacity traits caused by the suspended aqueous solutions. However, this cannot be counted on, and these traits can be drastically different for different reagents. The algorithm created to solve this problem has therefore been chosen to be both color-invariant and opacity-invariant, which is why work is being done with transparent, colorless droplets.

1.3 Previous Work

A great deal of research recently has been put into both gathering and interpreting feedback information from digital microfluidics systems, spanning many different methods. In particular, the work published by Shin[4, 5] in 2008 and 2010 was found to most closely parallel the work presented in the current study. Additionally, there can be found several papers by various authors[1, 7] dealing with using capacitance to check the state of the microfluidics system and additional papers by other authors[6] concerning using optical (or other) detectors on single electrodes at key points. However, all of these solutions have their problems.

1.3.1 Detectors in Literature

Using detectors (of any kind) at some few key electrodes, while probably a valuable idea for identifying the substance above the individual detector, cannot possibly detect the state of the entire system. Each detector can only say whether the system has performed correctly at the single point for which it has information. It cannot tell when, where, or more importantly, why the system malfunctions and cannot possibly hope to send a command to fix the mistake if the mistake could have happened at any one of a number of steps. To solve this problem, one could place detectors on most or all electrodes, but this would unnecessarily increase both the price and the complexity of manufacturing the device. It would also most likely make the system more susceptible to wearing out from fouling due to use.

1.3.2 Capacitance in Literature

Capacitance reading, as stated above, has both advantages and disadvantages for use as a feedback signal for a digital microfluidics device. Capacitance results in clear and easily understandable signals, but it cannot sense the current state of an electrode without first interrupting the voltage to the electrode in order to connect a capacitance reader to the electrode, thereby changing the state of the electrode. This is not to say that capacitance reading is not a valuable and potentially powerful solution to this dilemma, but while the problems that capacitance introduces into the system may be interesting to an electrical engineer, they are not computational problems to be solved with an algorithm in either the controller or the observer.

1.3.3 Vision in Literature

Other work involving computer vision solutions to microfluidics, as shown in Shin's papers,[4, 5] tend to be more concerned with solving the problem for specific applications such as glucose enzymatic assays, in which color and transparency can both be counted on as features that will stay constant throughout the use of the algorithm. In the current study, no assumptions are made about the properties of the reagents being handled on the digital microfluidics system. Instead, this thesis outlines the problems inherent in using classical computer vision algorithms on a digital microfluidics device while making few assumptions about the nature of the aqueous solution being utilized on the device and provides a possible solution.

2. THE ALGORITHM

2.1 The Basics

Computer vision is a broad, well-developed field with a very large metaphorical toolbox. There are algorithms for everything from image mosaicing to letter detection and differentiation to facial and emotion recognition. With a breadth as great as this, how does one determine the algorithm that best fits the application? Beginning with the options outlined in the literature, one of the most common algorithms for determining the location of an object is image subtraction, which involves a "before" image without the object in question and an "after" image with the object. When two images are subtracted from each other, this results in a third image with the area that differed between the two images being highlighted. That area corresponds to the object being measured. Unfortunately, image subtraction does not work with all digital microfluidics setups. In particular, on top-plate microfluidics systems, which suspend the ground above the surface, the top-plate typically must be removed to add the reagents, a process which can change the orientation of the device and the size, shape, and location of the silicone oil enough to cause the "before" image to be drastically different than any images gathered thereafter. Other major object detection algorithms in literature include the Scale Invariant Feature Transform (SIFT) and other feature-based techniques, which require the object in question to have constant distinguishing characteristics, a constraint that cannot be guaranteed with a transparent object, and Blob detectors, which require solid boundaries, something that can be hoped for but not counted on. After much deliberation, a Hough transform specialized for circle and ellipse detection was used. Before using the Hough transform, however, some steps had to be taken to preprocess the data.

2.2 Preprocessing

Some basic preprocessing steps were preformed on the collected images both to reduce the work of the algorithm and to help keep down the number of false results. First, since the algorithm is supposed to be color invariant, a grayscale effect was applied to the image, effectively cutting the image's complexity by half and reducing the size of each image file from 54kb to 29kb, a very noticeable jump with the quantity of images that were being processed. Next, the images were cropped down to contain only the area of interest instead of the entirety of the device and its surroundings.

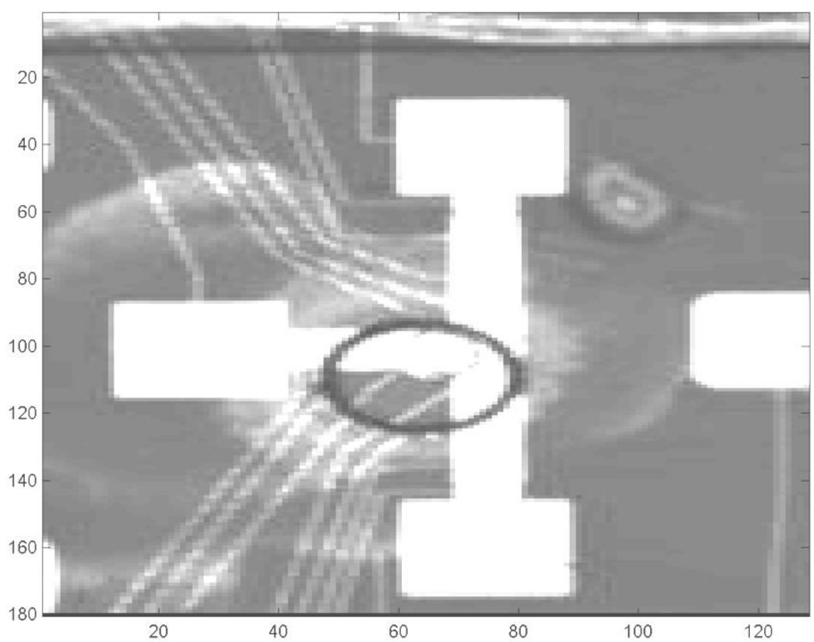

Figure 4: Grayscaled Image

2.3 Noise Removal

The next step in the algorithm is to apply a mask to the image to filter the results based on location and, sometimes, intensity. Clearly, since the intended result of the algorithm is to know the state of the device with respect to the electrodes, it is safe to assume that any circles distant from electrodes cannot be circles that are of interest. Therefore, since such circles are uninteresting, finding and classifying these circles is a waste of time for the algorithm. To handle this problem, a Canny edge detection function[2] was first performed on the image, but the original image was retained for later use in the algorithm.

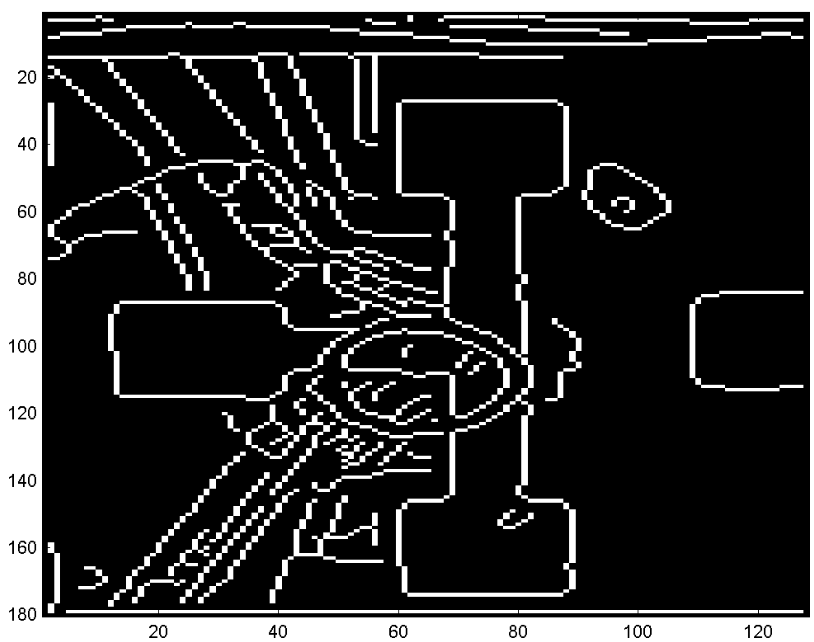

Figure 5: Canny Image

After the Canny Image had been created, an empty matrix of the same size as the image was created and a value of 1 was assigned to every pixel that was in close proximity to a white pixel (edge component) in the Canny Image. This gives a matrix that, if interpreted as an image, would look similar to the Canny Image, but with much thicker lines.

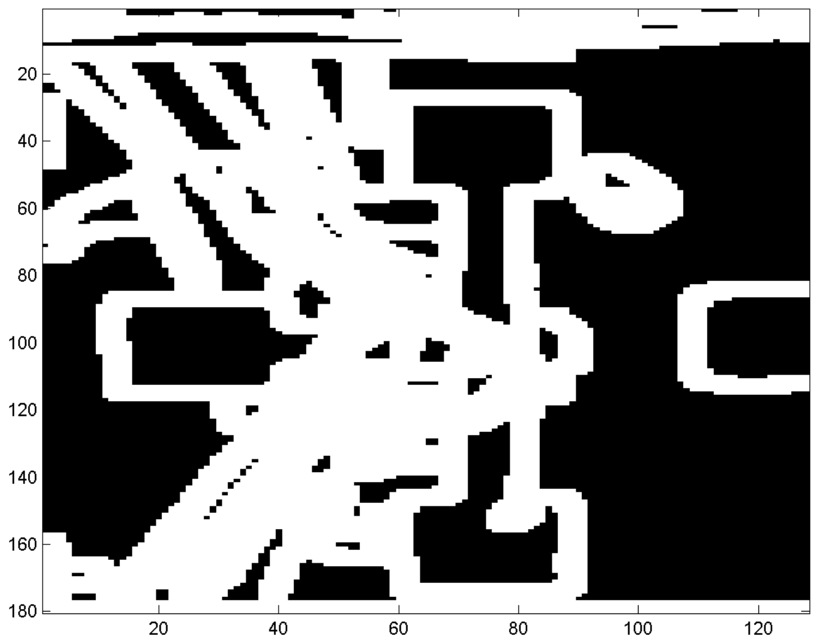

Figure 6: Canny Mask

Depending on system lighting, it is often effective to perform another similar operation based on intensity. In high-light situations, this is not particularly helpful, but in systems without proper lighting, it tends to give better results. Again, the original image is taken and this time a threshold operation is performed on it at approximately 10% more than the optimal contrast value as calculated using Otsu's threshold method[8] to obtain the brightest areas of the image. Again, an empty matrix of the same size as the image is created, assigning a 1 to pixels close to those remaining in the Otsu threshold image and a 0 to all others.

The resulting "mask" matrices are then multiplied by the original image, resulting in a masked image. In this masked image, pixels from the original image only survive if they are within a few pixels of an edge and within a few pixels of a bright area (if the intensity mask operation was performed), in which case they retain their original shade. Otherwise, they are assigned a 0 value, which interprets to black. After this step, however, another further step needs to be performed. The areas that were assigned a 0 in the previous step now form the most obvious edges where they abut the other pixels in the image. To fix this, the average pixel color on the original image is calculated, replacing all 0 values in the masked image with that average, thereby normalizing the image and causing the 0-value pixels to blend with the background. This result looks like a washed out version of the original image with clarity retained only in the areas near both brightness and edges, which now can have some more advanced image processing techniques performed on it without as many false positive results.

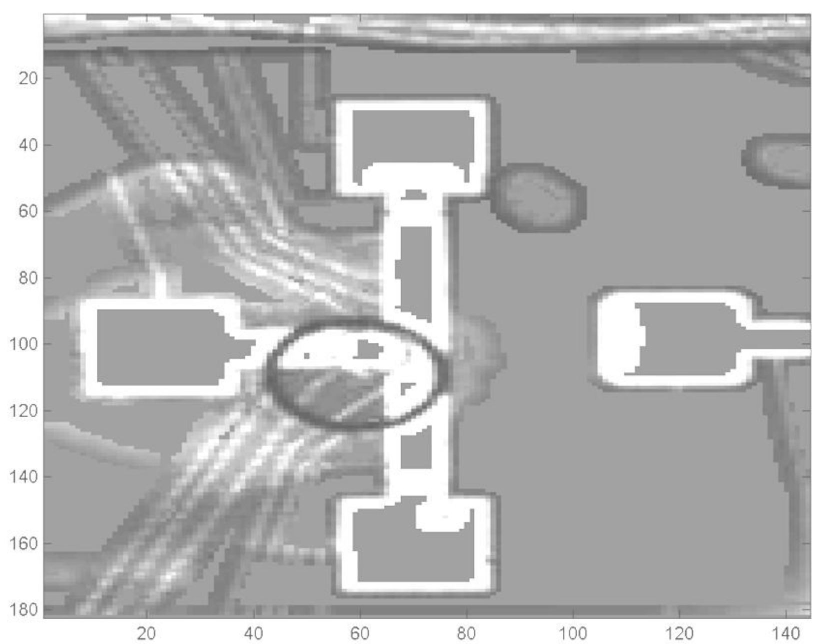

Figure 7: Masked Image

2.4 Hough Circle Detection

Next, the circle detection step is performed. Taking the masked and normalized image, each pixel in the image is examined. Based on the pixel's neighbors and the inherent gradient properties of edges, it is determined what circular edges this pixel could be a part of. A 'bin' is then created for each of these circles if such a bin did not exist already, defining them by their radius and the position of their center, and a 'vote' is placed in each of those bins. After evaluating this function for each pixel in the image, all of the votes in each bin are tallied and the circles that obtained the most votes are the most likely circles to be found in this image and, therefore, the locations of the droplets. If interpreted as an image, this set of bins, or Hough Accumulator, looks like a pattern of lights where brightness correlates to likelihood to be the center of a circle.

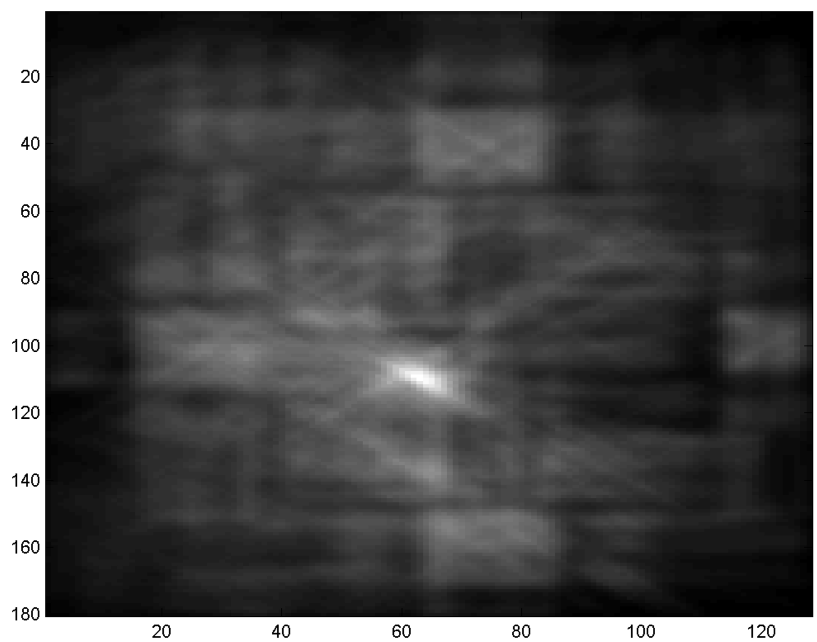

Figure 8: Hough Accumulator

The specific implementation used in this thesis is CircularHough_Grd, an implementation developed by Tao Peng[9] and modified by the author. CircularHough_Grd was chosen due to its well-documented code, clear variables, and speed. Although speed is still an important consideration, CircularHough_Grd performs the Hough Circle Detection algorithm in significantly less time than the other implementations that were tried, though additional work could be put into further optimization for speed, as the majority of the program's running time is still consumed by CircularHough_Grd. However, one notable disadvantage that should be brought to light is that CircularHough_Grd is only capable of detecting circles, not both circles and ellipses and, due to this, will often characterize an elongated elliptical droplet as two separate circular droplets. Since elongated elliptical droplets only tend to appear if a droplet is being split, this disadvantage is not a major one and is a suitable price to pay for a relatively low running time.

2.5 Adaptation to Video or Live Streaming

2.5.1 Sampling

Some changes had to be made to the code to adapt it from single images to video. First, the video had to be depacked into a series of images. Since videos tend to have multiple frames per second and the algorithm does not handle images nearly fast enough to take all of them in any reasonable amount of time, only a small subset of the frames were actually sampled. Sampling was found to work well with values ranging between 1/10th and 1/50th of the frames. A decline in success rate was not noticed anywhere over this range and because of this most tests only sampled one frame out of every twenty-five.

2.5.2 Noise Removal over Time

Adapting the algorithm to take video as an input also improved noise removal capabilities. It was often found that, on single images, some small false positives were seen to appear on many images with no seeming cause. When video was introduced, it was found that these were caused by changes in lighting such as shadows from movement and occasional bubbles in the reagent or oil. Since these were all very temporary, a noise suppression technique was utilized to remove them. For each circle that the Hough Circle Detector chose, that circle was only plotted if the same circle (or a similar one) had been found in the previous sampling. Because these samplings were often multiple seconds apart, this technique usually removed the false positives.

2.5.3 Live Streaming

Since the intended use of this algorithm is as an observer to detect feedback signals in the digital microfluidics device, much thought was put into adapting the algorithm for live streaming. The sampling choices made for this algorithm were necessary to obtain a reasonable running time, but this running time is not optimal for actual use on a digital microfluidics device. For live streaming, the algorithm would have to sample a frame from the stream, analyze it, and output its results in about one second. This would give commands a turnaround time of slightly more than a second, a reasonable timestep for this application. To optimize the algorithm to that level, however, one of two steps should be taken. Either the Hough transform must be greatly optimized or the algorithm should be run in parallel. Consider a multi-threaded system running this algorithm on each thread. The first thread would sample the data stream and recede to memory to process. A few frames later, the second thread could sample the stream and recede to memory as well. Once the first thread has finished processing, it can return to the data stream and sample again. The biggest difficulty for this implementation, however, is to find the optimal number of threads. Too many threads would lead to an incredible number of frames being processed but, since the controller would have to wait for each thread to finish processing before it can know whether the command completed successfully, there's not necessarily much speedup. To determine the optimal number of threads, it is necessary to break the experiment down into as many discrete modules as possible. If there are two commands being sent that cannot interact with each other, it is possible to send the first and have a thread processing the data from the stream while the second command is being sent. Clearly, this will not always be possible, but it could result in a great deal of speedup, potentially enough to allow this algorithm to perform in real time. However, not all solutions to these problems can be solved with new algorithms. Some require specialized equipment, as can be seen in the next section.

3. HARDWARE SOLUTIONS

3.1 Lighting

3.1.1 Difficulties with Lighting

Appropriate lighting is absolutely essential for any vision application and digital microfluidics is no exception. In fact, with digital microfluidics, proper lighting is especially important for inspecting transparent and reflective reagents such as water. Additionally, care must be taken to choose a lighting setup that is not hindered by the grounding slide, typically made of glass, which is suspended above the surface of the device. Several lighting solutions exist for this scenario, but they all have their advantages and disadvantages.

3.1.2 Backlighting

One potential solution which is typically not considered in literature is backlighting. In this situation, some form of light is placed under the digital microfluidics device and shines up through the surface of the device and into the camera. This technique, however, is at odds with the current practice for creating digital microfluidics devices, which calls for the electrodes to be made by starting with pre-cleaned glass slides coated with chromium (10 nm) and then gold (100 nm) by electron beam deposition and then etching off the gold and then chromium by application of a photomask, [10] a process that leaves very thin (but entirely opaque) electrodes on the surface of the device. Because of this, any backlight would be unable to properly illuminate the electrodes and, due to this, would not properly accentuate the edges of the droplets on top of the electrodes. Due to this flaw, backlighting is not considered a good lighting solution for droplet detection in digital microfluidics. Other branches of microfluidics, however, do not necessarily have this restriction and, therefore, may find backlighting to be the optimal solution.

3.1.3 Sidelighting

Another possible and much more useful lighting solution is sidelighting. With sidelighting, some form of light is placed to one side of the device and is meant to shine light across it, causing droplets to cast shadows to identify their location. However, as most implementations of digital microfluidics have a system for applying voltage to the electrodes directly around the device, the light source would need to be on the device to be unobstructed. Even then, the oil poses a problem. Since the droplet is surrounded by silicone oil in most implementations, instead of having the droplet cast a shadow, the oil would cast a shadow. An oil-less system, which some digital microfluidics devices employ, might find sidelighting to be an optimal solution but, in systems that employ silicone oil to aid movement, it is not considered to solve the problem. Despite that, some systems use sidelighting in other methods such as aiding the system in identifying the reagents contained in a droplet.[5]

3.1.4 Toplighting

A third solution for properly lighting a digital microfluidics device is to suspend the light above the device. This allows the light to shine onto the surface, reflect off the droplets, and travel back to the camera. The difficulty with using toplighting, however, is choosing the correct type of light to use and the location to place it. If either of these two choices is made incorrectly, a variety of problems can occur. Instead of reflecting off the droplets, the light could reflect off the surface below them or the grounding plate above them. Instead of highlighting the small edges of the droplet, the light could highlight the oil or the minute bubbles and flaws in the device that occur naturally through use.

After speaking with many experts, it was determined that an Advanced Illumination Extended Working Distance LED Ring Light from Edmund Optics (part no. 66784) was an optimal choice for this application. For other individual implementations, it is recommended that the experimenter contact a specialist and try several different options. As can be seen, the LED Ring Light causes much more even highlighting of the droplet than typical laboratory lighting.

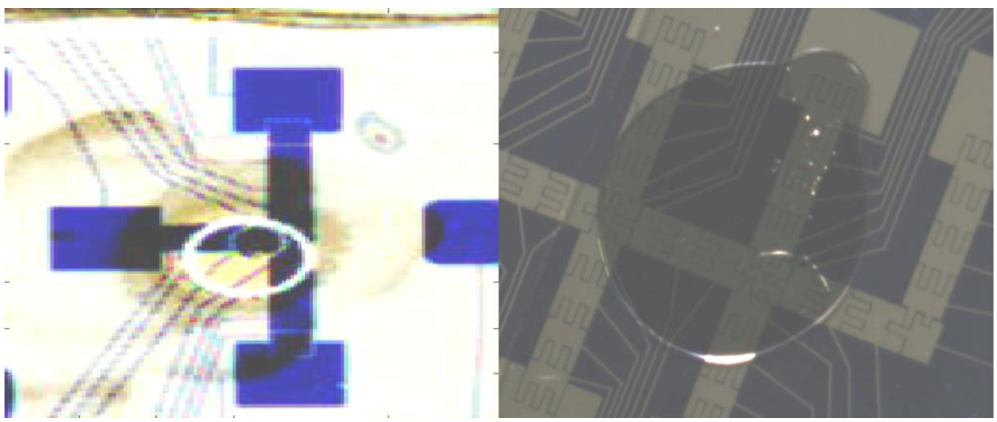

Figure 9: Top Light vs. Side Light

3.2 Oil Gasket

One of the most difficult problems with vision and digital microfluidics is the misclassification. If the algorithm is looking for circular shapes in an image which has a large, obvious oil pool and a much smaller, less obvious water droplet, the oil droplet is much more likely to be found. An oil gasket is a piece of material that sticks to the surface of the microfluidics device outside of view of the camera. If the gasket is constructed and applied correctly, it allows the oil to fill the entire viewing area without the danger of it leaking out and damaging the equipment running the device. Using an oil gasket removes the problem of misclassifying the oil pool as an object of interest, thereby simplifying the problem. Through various tests, it was found that PDMS (Polydimethylsiloxane) is quite capable as a material for creating such a gasket, as it adheres well to the top layer of the surface of the digital microfluidics device with an oil-tight seal.

3.3 Oil Dyes

Another possible hardware solution is to obtain an oil-soluble dye that is not water-soluble and dye the oil a specific color (black would likely work best with the current implementation). This would both allow the addition of a color-dependant step to the algorithm to discount all parts of the image that were the same color as the oil. Clearly, the color of the dye would have to be chosen carefully so as to be easy to differentiate from all of the reagents, the addition of such a step could drastically reduce the processing time of the algorithm. Additionally, colored oil would also highlight the contrast between the oil and the droplet, thereby increasing performance. However, a suitable oil-soluble, non-water-soluble dye could not be determined.

3.4 Single-Plate Devices

This final hardware solution regards making a large change with the device. While many digital microfluidics devices suspend the ground plate over the surface of the device, some devices, known as single-plate microfluidics devices, ground the droplet from below by using thin conductive lines on top of the insulator layer. Since the device no longer needs to be grounded from above, some of the difficulties inherent in finding the correct lighting solution are no longer relevant. Additionally, it may be possible to perform operations on a single-plate device without the use of oil, thereby removing further problems. However, some of the drawbacks to single-plate microfluidics make it unsuitable for certain applications.

In particular, single-plate microfluidics devices are much more susceptible to evaporation than top-plate devices and may not be capable of splitting droplets.[11]

4. RESULTS AND DISCUSSION

4.1 Data Collection

The data for this thesis was collected over several years, working on a top-plate digital microfluidics device developed at Linhardt Labs at Rensselaer Polytechnic Institute. Attempts were made to collect data with different lighting setups and this algorithm was tested on both side-lit images and top-lit images. Images with no lighting solutions did not have enough contrast to determine the locations of the droplets and back-lit images were not feasible with our setup.

4.2 Side-lit Results

The results gathered with sidelighting had several large disadvantages. Because the light was mostly directed across the surface of the device, many of the images were dark with only a small amount of highlight on some of the edges of the droplet, as can be seen in Figure 10.

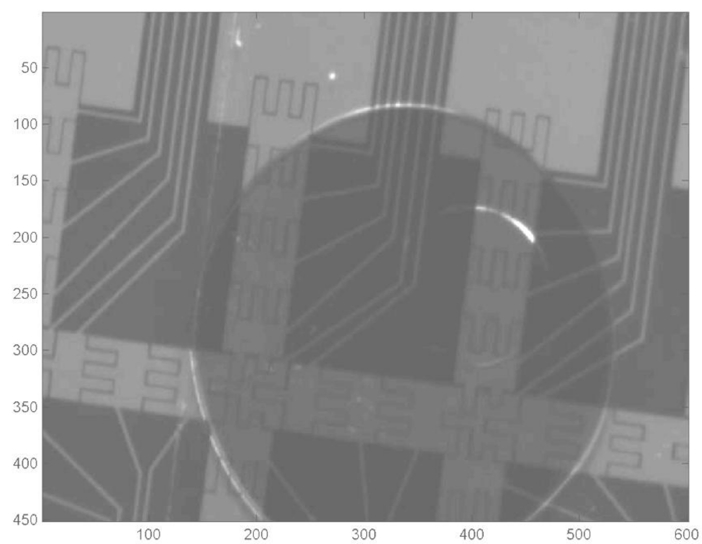

Figure 10: Sidelighting Data

Additionally, a similar highlight was found on some of the edges of the oil pool, causing the algorithm to confuse the two. It was found that adding the optional intensity mask step typically improved the results, but occasionally introduced additional false-positives. As can be seen in Figures 11 and 12, the intensity mask step removed a great deal of unnecessary detail inside the oil droplet, and it is hypothesized that the addition of a oil gasket would cause this to occur over the entire surface, removing the unnecessary background detail.

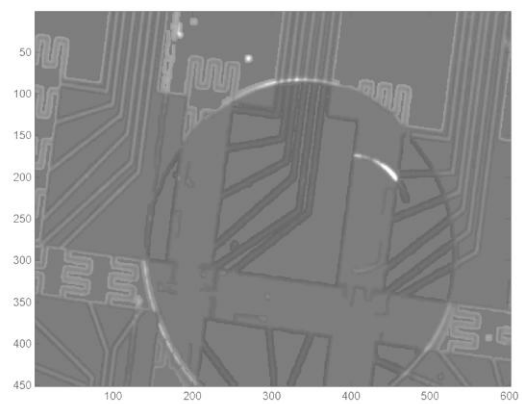

Figure 11: Sidelight Not Intensity Masked

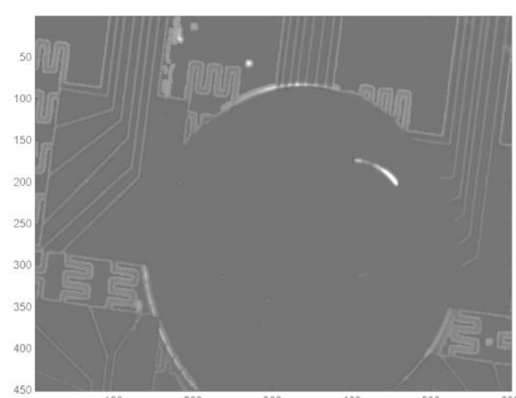

Figure 12: Sidelight Intensity Masked

However, though unnecessary detail was successfully removed from the image by the intensity mask, the droplet edge is still quite weak and difficult to identify as a circular edge. It is possible that the algorithm would be more successful if the oil were dyed, thereby further outlining the droplet edge. Figure 13 shows the circles detected by the algorithm on the intensity masked image. As can be seen, the droplet is identified, and several false positives are also present. Most of the false positives identified by the algorithm, however, are on the edge of the oil pool. Because of this, many of these false positives are likely to be eliminated by an oil gasket. The biggest problem with this image is that the droplet only has one highlighted edge: the leading edge in the direction of the light. Due to this downside of the lighting solution, the algorithm is not able to correctly identify the center or the radius of the droplet and is only able to make an estimate. Though it is difficult to see, the marked circle is too small and should be about ¼ bigger.

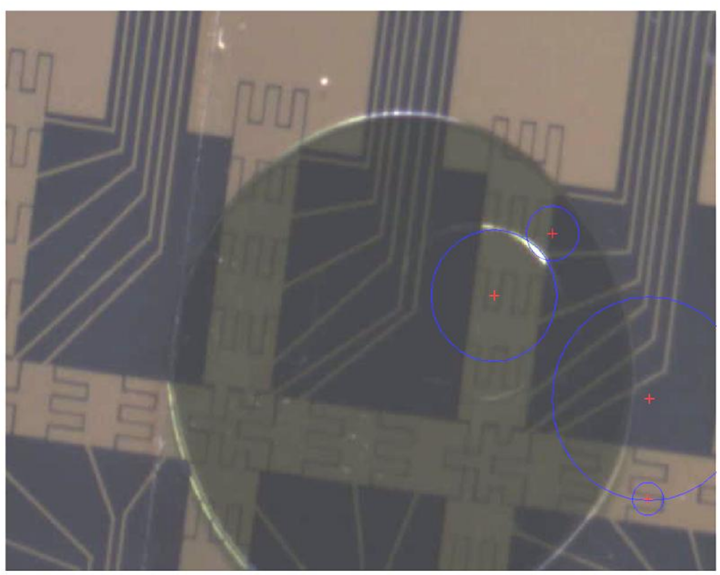

Figure 13: Sidelight Circles Detected

4.3 Top-lit Results

Since the side-lit results had such noticeable failings, several top-lit results were also obtained and tested with the algorithm. The images were found to be much brighter than the side-lit results, which was to be expected. Something unexpected, however, was that the outline of the droplet was dark rather than bright. This is believed to be because the curvature of the edges of the droplet cause the light to be reflected in a different direction, leading to light returning to the camera nearly everywhere else in the image but not at the droplet's edges. Because of this, since the edges are not necessarily the brightest parts of the image with top-lighting, the intensity masking step was determined to produce poor results as can be seen in Figures 14 and 15 and was not performed. A lot of unnecessary detail is removed through intensity masking, but so is a quarter of the droplet's edge, detail that we would like to retain.

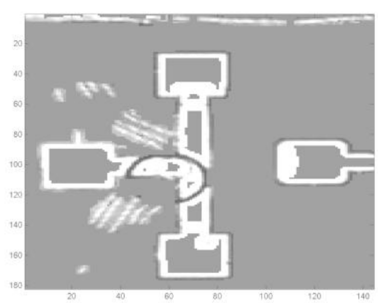

Figure 14: Toplight Intensity Masked

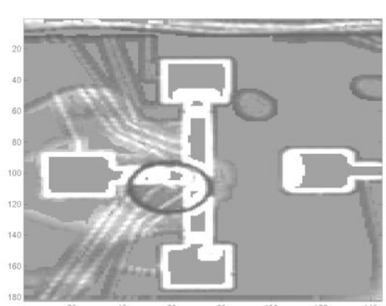

Figure 15: Toplight Not Intensity Masked

In this case, contrary to the sidelighting case, the edge can be seen to be strongly highlighted around the entirety of the droplet. After the Hough voting step is performed, we are given a very well defined Hough Accumulator as shown in Figure 16 (below).

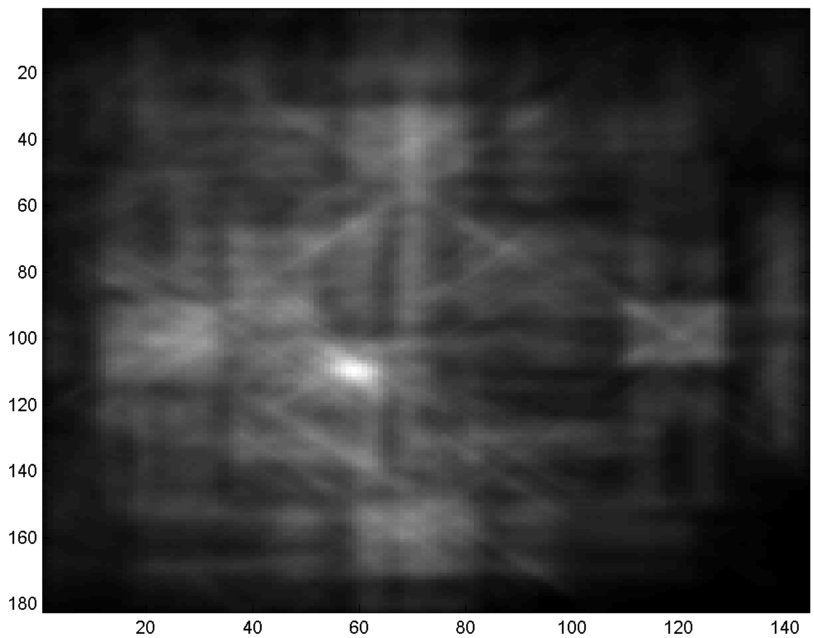

Figure 16: Toplighting Hough Accumulator

As can be seen in the accumulator, the droplet itself obtains the vast majority of the votes, though some other locations are also voted on. In particular, the large square reservoir electrodes on the edges of the also show up noticeably in the accumulator. As shown in Figure 17, these votes are even enough for several of the square reservoir electrodes to be chosen as circles.

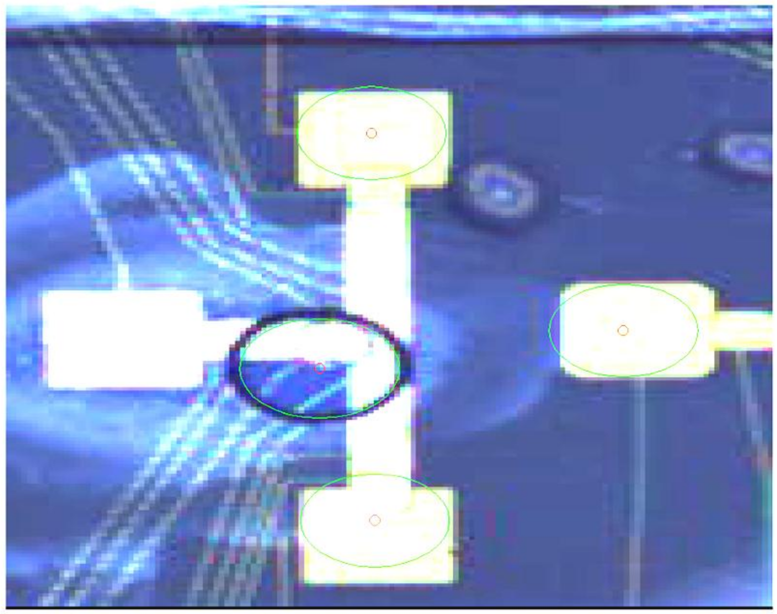

Figure 17: Toplighting Circles

This was not expected, but is not much of a disadvantage for two reasons. Firstly, in actual operation, the reservoirs will actually have droplets on them at all times and, therefore, should be selected as locations which contain droplets. Secondly, droplets that are part of reactions (the ones that need to be tracked) will never be on these reservoirs. Finally, these circles will be less likely to appear if the reservoirs are properly covered with oil as the reservoir on the left is.

4.4 Conclusions

Although the algorithm does not give perfect solutions to the problem, this thesis has shown that a modified Hough Circle Detection algorithm has high performance for the application of droplet detection on a digital microfluidics device. In addition, it has been shown that, with slight variations, such an algorithm both performs well in a situation with poor lighting and shows noted improvement in a situation where the lighting has been selectively tailored to the specific digital microfluidics system setup. Specifically, this thesis shows that, in a low light condition in which at least one of the droplet edges has still been highlighted by sidelighting, the algorithm is likely to misjudge the actual size of the droplet, whereas in a well-tailored lighting condition where all of the droplet edges are highlighted by toplighting, the algorithm does not have such difficulties. Also, the algorithm has been shown in both cases to find several false positives in areas nearby the edge of the oil pool surrounding the droplets, a problem which a solution is suggested to rectify.

4.5 Future Work

In conclusion, several advances have been made in addressing the problem that this thesis sought to solve, but many additional questions have arisen. Several solutions that could be implemented by digital microfluidics systems to improve performance of vision-based observer algorithms for the purposes of droplet detection are postulated in this thesis. These solutions should improve the accuracy and/or speed of the modified Hough Circle Detector algorithm (or many other algorithms) if they were to be implemented. Clearly, more research needs to be done to determine whether these solutions will be as successful as postulated. In specific, more testing should be done to verify the plausibility of an oil gasket made of Polydimethylsiloxane in order to fill the entirety of the viewing area with silicone oil without risking damage to the running equipment and the digital microfluidics device itself. In addition, it is necessary to look further into the possibility of oil-soluble colored dyes that are not water-soluble, do not react with the reagents stored within the droplets, and would give increased contrast if compared with the reagents of the particular reaction. Finally, though the modified Hough Circle Detection algorithm does not operate at optimal speed to interface with a live input stream in its current form, suggestions have been made detailing modifications to the current algorithm to increase efficiency and allow parallelization for further speedup. Clearly, a system must be devised which tests the parallelization of this algorithm. Currently, the system fulfils expectations, finds and identifies droplets, and succeeds with good performance in the general case. The problem which was presented here is only interesting in such a general layout to a point and from here should be specialized to fulfill the needs of each specific system. This process may provide new, interesting problems and solutions to more detailed and constrained forms of the problem described in this thesis.

BIBLIOGRAPHY

1. Martin, Jeffrey (2010) *Creation of an Artificial Golgi Organelle: Glycosaminoglycans in Microfluidics.* , unpublished thesis (Phd), Rensselaer Polytechnic Institute
2. Canny, John (1986) *A Computational Approach to Edge Detection*, Transactions on Pattern Analysis and Machine Intelligence, Vol 8(6), 679-698
3. Kierkegaard, P. (1992) *A Method for Detection of Circular Arcs Based on the Hough Transform*, Machine Vision and Applications, Vol 5, 249-263
4. Shin, Y. and Lee, J. (2010). *Machine vision for digital microfluidics*, Review of Scientific Instruments, Vol. 81
5. Shin, Yong-Jun (2008) *EWOD-Based Digital Microfluidics as a Novel Platform for Systems Biology*, unpublished thesis (MS), University of Texas at Dallas
6. Fair, R. (2007) *Digital microfluidics: is a true lab-on-a-chip possible.* Microfluid Nanofluid, Vol 3, 245-281
7. Chen et al. (2004) *Capacitive sensing of droplets for microfluidic devices based on thermocapillary actuation.* Miniaturisation for Chemistry, Biology & Bioengineering, Vol 4, 473-480
8. Otsu, N (1979) *A Threshold Selection Method from Gray-Level Histograms*, IEEE Transactions on Systems, Man, and Cybernetics, Vol. 9(1), 62-66
9. Peng, Tao (2005) *MATLAB Central - File Detail - Detect Circles with Various Radii in Grayscale Image via Hough Transform.* http://www.mathworks.com/matlabcentral/fileexchange/9168 date last accessed Apr. 2011
10. Watson et al. (2006) *Microcontact Printing-Based Fabrication of Digital Microfluidic Devices*, Anal. Chem, Vol 78(22), 7877-7885
11. Cooney et al. (2006) *Electrowetting droplet microfluidics on a single planar surface*, Microfluid Nanofluid, Vol 2(5), 435-446