Fast Reactive Control for Illumination Through Rain and Snow

Raoul de Charette[1], Robert Tamburo, Peter C. Barnum[2],
Anthony Rowe, Takeo Kanade, Srinivasa G. Narasimhan
The Robotics Institute, Carnegie Mellon University, USA
[1]Mines ParisTech, France
[2]Texas Instruments, Inc., USA
http://www.cs.cmu.edu/~ILIM

Abstract

During low-light conditions, drivers rely mainly on headlights to improve visibility. But in the presence of rain and snow, headlights can paradoxically reduce visibility due to light reflected off of precipitation back towards the driver. Precipitation also scatters light across a wide range of angles that disrupts the vision of drivers in oncoming vehicles. In contrast to recent computer vision methods that digitally remove rain and snow streaks from captured images, we present a system that will directly improve driver visibility by controlling illumination in response to detected precipitation. The motion of precipitation is tracked and only the space around particles is illuminated using fast dynamic control. Using a physics-based simulator, we show how such a system would perform under a variety of weather conditions. We build and evaluate a proof-of-concept system that can avoid water drops generated in the laboratory.

1. Introduction

Imagine driving at night in a rainstorm. The bright flickering rain streaks, the splashes and the blurry windshield (Fig. 1) all contribute to the degradation of scene visibility [12], making driving an unpleasant experience at best and a dangerous endeavor at worst. In this paper, we present a fast imaging and illumination system for reducing both light reflected back to the driver as well as light scattered towards nearby and oncoming drivers.

Recent work in computer vision proposes to digitally process images to remove the effects of fog, haze, rain and snow [17, 9, 18, 2, 21]. These methods can be used to initiate driver-assistance tasks. However, the resulting weather-free videos need to be shown on a heads up display for the driver. Current implementations of these types of systems are unnatural and often distract drivers. We propose an approach that directly removes the appearance of rain from the scene without any disruption to the driver. The approach takes advantage of the fact that rain only becomes visible when it is illuminated by various light sources on the road.

Figure 1. At night, illuminated rain or snow appears as a bright flickering (distracting) pattern that reduces driver visibility (top). We propose a reactive system (bottom) that deactivates only those light rays that intersect particles such as raindrops, snowflakes, and hailstones diminishing the visual effect of precipitation.

A rain avoidance system would need to first identify where the raindrops are located and then reactively deactivate rays of light to avoid illuminating them. There are significant challenges in building such a system. First, the system response must be fast enough to avoid precipitation particles (i.e., raindrops, snowflakes, and hail). Existing work on anti-blinding headlights [4, 13] use adaptive illumination to avoid glare for the drivers of oncoming vehicles, but are too slow to be suitable for our application[1]. Second, the system must maintain a high level of *light throughput* so that the scene remains sufficiently illuminated. This requires both a low-latency high-speed system as well as accurate tracking and prediction algorithms.

In this paper, we provide both simulation and a prototype that suggest such a system is possible. There are three main observations that help support the feasibility of our system. First, due to light fall-off, it is not necessary to consider

---

[1]A $3mm$ raindrop falling at a few meters per second is equivalent to a large ($3m$) car traveling at several thousand kilometers per hour.
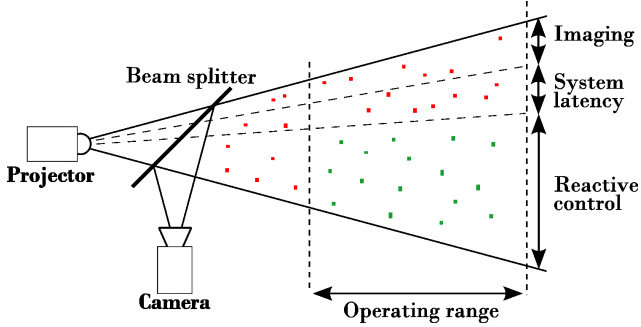
Figure 2. A co-located camera-projector system images and illuminates a volume of precipitation. All particles are first detected by illuminating them in a very short period of time, their future locations are predicted, and then the rays intersecting them are reactively turned off. The time duration between image capture and reactive illumination is the latency of the system. The operating volume is approximately $3m \times 3m \times 4m$.

Figure 3. Pipeline stages of our system with execution times (in ms). TX_c and TX_p denote data transfer between camera and computer and between computer and projector, respectively. Process refers to drop detection and prediction, and generation of the projection image.

drops that are beyond a relatively short distance from the light source. Second, it is not necessary to consider drops that are too small to reflect a significant amount of light [15]. Finally, strong visibility improvement can be achieved well below perfect accuracy of the system. For example, by avoiding 70% of the drops near the headlight while losing 5 to 10 percent of the light still improves the driver's visibility.

We demonstrate using quantitative simulations that the light throughput of an ideal system suffers very little for a broad range of precipitation types and rates. Based on the lessons learned from the simulations, we present a prototype system and describe the trade-offs between the algorithm used, hardware speed, and achievable light throughput and accuracy that are needed to make the system realizable as a vehicular headlight.

2. Overall Approach

The components of our co-located imaging and illumination system are shown in Fig. 2. At the top of the camera's field of view (FOV), particles (raindrops, snowflakes, and hailstones) are illuminated by a projector and observed within a short period of time (typically, a few milliseconds). This region of the image is used to predict the trajectory of particles across following frames. In the tracking region, light rays are prevented from illuminating particles by projecting around predicted particle locations. Simultaneously, any wrongly illuminated particles are detected by the camera and used to update the predicted locations. Co-locating the camera and projector simplifies tracking since the positions of particles only need to be determined in a 2D image instead of the full 3D coordinate space.

We define *system latency* to be the time required to transfer an image, process the data, and reactively control an illumination source. A system that suffers from long latencies will perform poorly for two reasons: (1) The particle may have already left the imaging/illumination FOV before the next control cycle, and (2) effects like wind and camera vibration could increase tracking error. In order to increase system responsiveness, many of the tasks can be executed in parallel, which improves processing throughput.

The timing diagram of our three-stage processing pipeline is shown in Fig. 3 with times measured from our prototype system (Section 4). Capture refers to camera integration and TX_c refers to image data transfer to the computer. Process refers to drop detection and prediction, and generation of the projection image. TX_p denotes the computer-to-projector transfer time and projection is the refresh time of the projector. The timing values show that the typical execution time of each component and the pipeline is staged based on the overall critical path (in this case stages 1 and 2 are almost identical).

We evaluate the performance of our system using two competing metrics. First, *light throughput* is a measure of the amount of light radiating from the light source. This is computed as the percentage of projector pixels that remain on per each frame. Second, *accuracy* is the percentage of particles that are not being illuminated. Thus, with a headlight that is constantly on, light throughput is 100% and accuracy is 0% and when the headlight is off the light throughput is 0% while the accuracy is 100%. Given these two competing metrics, our goal is to maintain high light throughput while maximizing accuracy.

Before implementing our system, we designed a simulation environment to explore the following design-space questions: What intensities of precipitation (drizzle, thunderstorm, blizzard, etc.) can the system handle? How much light does the system lose? How far from the light source or viewer should the particles be considered? How will the speed of the vehicle effect performance? How fast should the reactive control be? What is the trade-off between system latency and complexity of the prediction algorithm?

We demonstrate a low-cost system built with off-the-shelf components that can achieve effective visibility improvement without losing significant light throughput. The simulations are described next.

3. Feasibility Study Using Simulations

Our discrete-time simulator is capable of emulating the sizes, densities and dynamics of particles typical to rain, snow, and hail. We generate particle size distributions, velocities and dynamics from the physics-based models in literature [16, 19, 14, 11, 6]. The simulator provides a virtual camera and projector to model the impact of additional system parameters including camera exposure time, camera resolution, processing delays, projection resolution, and projection refresh times. The components in the simulation are updated in an asynchronous manner so as to accurately capture timing dependencies. At the end of this section, we report simulation results and provide concrete recipes for an actual system.

3.1. Particles System Simulator

The density and number of particles are directly proportional to the precipitation rate R (mm/hr), which is defined as the amount of water that falls on the ground per hour. If the area on the ground where particles fall is A (m^2), then $A * R$ is the total number of liters collected in 1 hour. As our simulator is updated with a frequency of P_σ (Hz) the amount of water W (L) to fall is computed for each time-step using:

$$W(t + \delta t) = W(t) + \frac{AR}{3600 P_\sigma}. \tag{1}$$

Particle size distributions (PSD) are modeled using exponential functions scaled by the precipitation rate R (i.e. large particles are more likely to occur in heavy precipitation). Given the weather parameters, we compute the probability function modeling the corresponding PSD and use it to determine the size of each upcoming particle. Each time a particle is generated, its *water equivalent* is removed from the remaining water and the same process is repeated. For snowflakes and hailstones, the water equivalent is the result of the melting process assuming a water density of 30% and 80%, respectively. For example, $10mm/hr$ water equivalent of snow in our simulator corresponds to $0.8m$ of accumulated snow in one day.

The sizes, PSDs, and velocity models used in the simulator are listed in table 1. In most natural scenarios, it is sufficient to model the dynamics of particles using a constant terminal velocity. If particles are artificially generated (say, using sprinklers), initial velocities and acceleration must also be taken into account [19]. We operate our
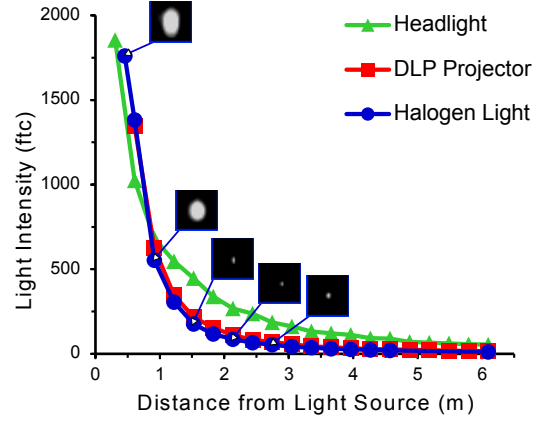

Figure 4. In the laboratory, $4mm$ water drops were illuminated by a halogen light and imaged for comparison at different distances. The plot shows that water drops are not visible farther than $3m$ with an exposure time of $4ms$. Also shown is the light fall-off for a Honda Civic (2006) headlight (late night with no street lamps) and Viewsonic PJD6251 DLP projector (color wheel removed and illuminated through a 50% beamsplitter).

simulator at a frequency of $P_\sigma = 10000 Hz$ providing a $0.1ms$ time granularity.

3.2. Operating Range

The area covered by the projector increases as the distance from the projector increases, which in turn increases the probability of hitting more particles. Fortunately, the depth of interest is limited by light fall-off and the small size of particles. In order to quantify this effect, we captured images of $4mm$ (diameter) water drops, illuminated by a halogen lamp at different distances from the camera. Figure 4 shows the light intensity with respect to distance given a 1024×1024 image and a $4ms$ exposure time. Further than $3m$, the droplets are no longer visible and thus the particles should cover $3m$ depth[2]. Additionally, width ($3m$) and altitude ($4m$) are set to cover the camera's FOV. This corresponds to a $3m \times 3m \times 4m$ volume within our simulator. Also shown in Fig. 4 is the light drop-off of a 2006 Honda Civic headlight (at night with no street lights present) and a Viewsonic PJD6251 DLP projector (color wheel removed and illuminated through a 50% beamsplitter).

The relationship between the integration time of the camera and the velocity of the particles is directly responsible for the image area covered by the particles [10]. Figure 5 shows rain and snow captured with a $30ms$ and $1ms$ exposure time. For each type of precipitation, longer exposure times create image streaks (similar to what a human perceives) and reduces the light throughput.

---

[2]This result was also confirmed by informal observations by four individuals in real rain and rain generated using sprinklers.

Table 1. Listing of the particle sizes, article size distribution, and velocity models used to simulate rain, snow, or hail. R is the precipitation rate in mm/hr, $N(d)$ is the number of particles of diameter d (mm) in 1 m^3. In many case, terminal velocity v_t (m/s) is sufficient to simulate real rain, snow or hail. To study the effect of initial velocity and wind drag, additional numerical models [19] must be used to compute the acceleration $a(t, d)$ (in m/s) as a function of time t and drop diameter d (omitted here for brevity).

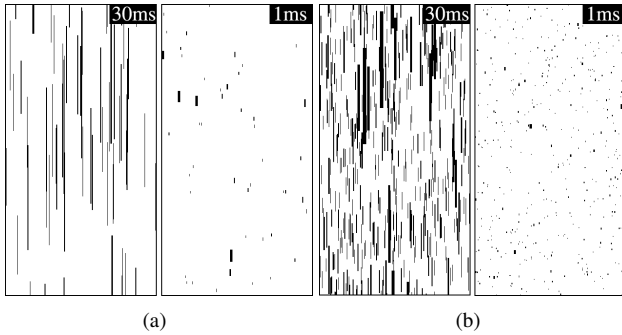|	Diameter	Particle Size Distribution $N(d) = N_0 * exp(-\lambda d)$	Velocity
Rain	0.1 - 10mm	$N_0 = 8000$, $\lambda = R^{-0.21}$ (source [16])	$v_t(d) = 9.65 - 10.3 * exp(-600 * d)$ (source [1])
Snow	1 - 10mm	$N_0 = 3800 * R^{-0.87}$, $\lambda = 25.5 * R^{-0.48}$ (source [11])	$v_t(d) = 2. * (10 * d)^{0.31}$ (source [14])
Hail	1 - 50mm	$N_0 = 12.1$, $\lambda = -0.42$ (source [6])	$v_t(d) = \sqrt{\frac{4*10^3*g*\rho_i*d}{3*\rho_a*0.6}}$ (source [5])

(a)   (b)

Figure 5. Pixel area covered by particles in moderate rain ($5mm/hr$) (a) and snow ($2mm/hr$) (b) when imaged with $30ms$ and $1ms$ exposure time. A longer exposure lowers light throughput. Thus, the camera exposure time should be no more than a few milliseconds to ensure meaningful performance.

3.3. Performance Analysis for a Stationary System

We begin by considering a theoretical system where image capture, processing, and projection are instantaneous and tracking is perfect. The accuracy is 100% and the light throughput can be used as an upper bound for what is possible in reality. Figure 6 shows the light throughput achieved for a variety of weather conditions including mild, moderate, and heavy rain, snow and hail. Notice that in the case of zero exposure time, the light throughput is above 99% for a thunderstorm ($17mm/hr$) and hailstorm ($10mm/hr$ equivalent quantity of water), and above 98% for heavy snow ($10mm/hr$ equivalent quantity of water). These results show that reactive illumination in bad weather could theoretically provide very high light-throughput. We now investigate system performance with parameters based on current technology.

The impact of camera exposure time on light throughput is shown in Figure 6. We see that to guarantee light throughput of 90% under most weather conditions, camera exposure time should be no longer than $2ms$. The time period for which particles stay in the projector's FOV also varies with different types of weather conditions (Figure 7).
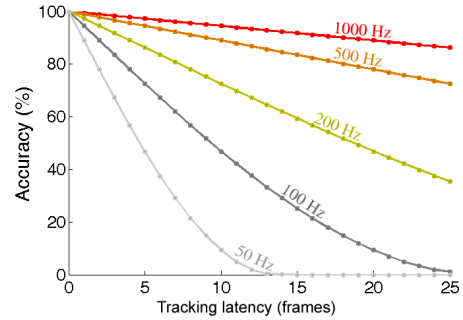

Figure 8. System accuracy versus tracking latency (measured in frames) for various update frequencies. Tracking may require several frames to initialize, lowering the accuracy of the system especially a with low update frequency.

Snowflakes tend to stay longer in the FOV than rain or hail. This has two effects — the light throughput will be lower in snow, but at the same time better accuracy may be achieved since a higher system latency can be accommodated.

We also investigate the impact of latency on tracking accuracy in our system. Assuming a raindrop diameter of $3mm$ and a response time of $10ms$, the drop will have moved by $80mm$ (26 times its own size) by the time the projector illuminates the scene. Figure 8 shows that less than 6% of the raindrops will be illuminated if tracking initialization requires 10 frames and the system operates at $1000Hz$. In contrast, at $100Hz$ more than 50% of the drops will be illuminated. Figure 9 demonstrates that light throughput decreases with increasing error in drop detection — an error of 5 pixels will result in 95% light throughput.

3.4. Performance Analysis for a Moving System

We extended our simulator to include a moving platform (vehicle) and assessed light throughput and accuracy for different vehicle speeds under various precipitation rates. We report results for a system operating at $120Hz$ with 3 frames latency and 2 pixel detection/tracking error. Figure 10 exhibits the influence of vehicle speed on accuracy in rain,
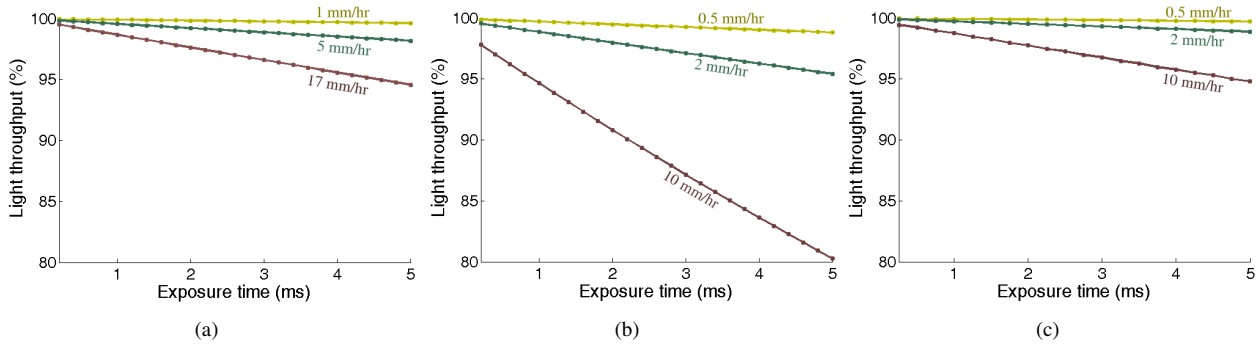
Figure 6. Ideal light throughput of our system for rain (a), snow (b) and hail (c) versus exposure time of the camera. Three precipitation rates are reported (light, moderate, and heavy). Particles falling slower stay longer in the field of view and thus require shorter exposure time. To achieve light throughput over 90% the exposure must be set to $2ms$ or less to accommodate the various precipitation types and intensities. For just rain and hail, $4ms$ is short enough.
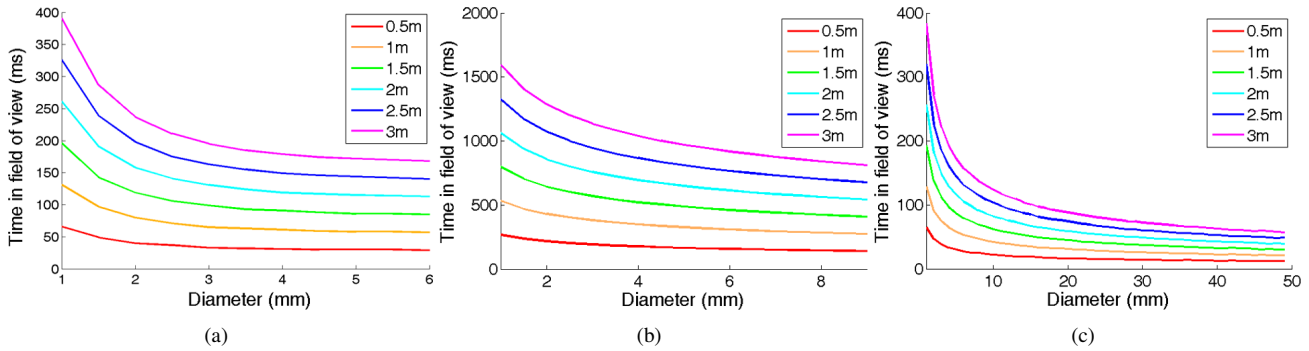

Figure 7. Duration particles of varying diameter stay in the field of view with our camera settings for rain (a), snow (b), and hail (c) given various distances from the camera. Avoiding illumination of raindrops or hailstones requires significantly faster response time since they fall faster than snowflakes.
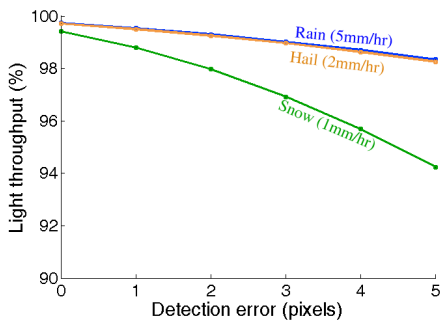

Figure 9. Light throughput versus detection error for rain, snow, and hail. As expected, throughput decreases when errors occur.
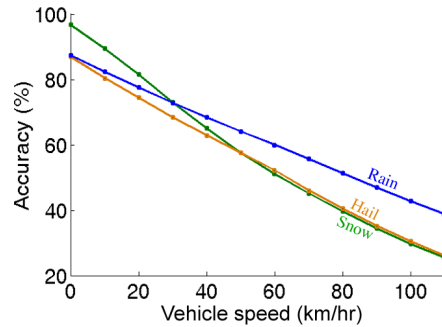

Figure 10. Accuracy versus vehicle speed in rain, snow, and hail. As the speed increases the accuracy severely drops but even at high speed our system successfully *not* illuminate some of the particles. Note that the accuracy is the percent of particles not illuminated which is not related to the density of particles (i.e. weather intensity). The accuracies displayed in this plot for rain, snow or hail stand for any weather intensities.

snow and hail. As the accuracy measures the percent of particles not being illuminated, it is not related to the weather intensity (i.e. particles density). For example, with a vehicle moving at $60km/hr$, the system accuracy for any intensity of rain, snow, or hail is respectively, 59.98%, 51.09% and 52.20%. Conversely, system accuracy is closely related to the duration for which the particles stay in the FOV (which is very short when the vehicle is moving fast).

Now, let us discuss the plot for snow in Figure 10. In contradiction to the good performance with a stationary system, snow accuracy drops dramatically. This is because the
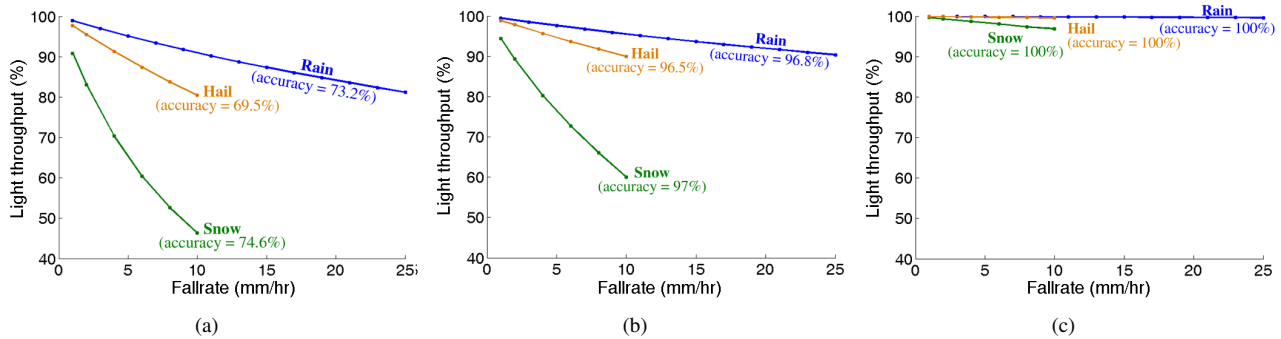
Figure 11. Simulation performance for three systems at $30km/hr$ using today's low cost (a), future (b) and ideal (c) hardware and software technology for various precipitation rates. All of these systems achieve acceptable accuracy ($> 69\%$). However, today's system (a) has $13ms$ latency and thus can't handle extreme conditions or snow without losing too much light. System (b) with $1.5ms$ latency gives us hope concerning the performance that can be achieved in near future. Not only the accuracy is over 96% but the light throughput is very high. The ideal system (c) is the upper-bound for what is possible with an instantaneous and perfect system. Notice that snow exhibits higher accuracy since particles fall slower although the throughput is significantly lower. Refer to section 3.5 for complete settings.

snowflakes fall slower than other particles and thus they remain longer in the air. Each time the vehicle moves forward a large number of flakes enter the FOV and will be illuminated before being tracked 3 frames later. However, notice that even at high speeds, our system reduces precipitation visibility. With a vehicle moving at $110km/hr$, 38.74% of the raindrops, 26.02% of the hailstones, and 25.48% of the snowflakes, in headlight's FOV, will *not* be illuminated.

In addition to vehicle motion, the effect wind on system performance should be investigated. We expect that wind does not pose a significant problem in many scenarios because particles remain within the field of view for a short duration of time (few milliseconds). It may be assumed that the wind direction mostly remains stable within this time period. However, if the wind is too chaotic, higher tracking latency will be required. Using the same system parameters described above with 6 frame latency instead of 3 frames, accuracy (and light throughput) in $5mm/hr$ rain is 74.68% (96.56%) for a stationary vehicle, is 51.00% (96.92%) for a vehicle moving at $30km/hr$, and is 13.74% (98.60%) for a vehicle traveling at $100km/hr$. Thus, in extreme situations, our system functions almost like a normal headlight with only a 14% reduction in precipitation visibility.

3.5. Recipe for a Complete System

We now present results at $30km/hr$ from two candidate systems with parameters that are possible to realize with currently available technology. In the first system, camera and projector resolution is 1024×1024 pixels, camera exposure is $1ms$, tracking latency of 3 frames, detection error of 2 pixels and with a camera motion of $30km/hr$. Based on USB-2 data transfer limitation, we assume our system to have a response time of $13ms$ (pipeline identical to Figure 3). The plot in Fig. 11(a) shows the results for different

precipitation rates of rain, snow and hail. Notice that the accuracy of the system is quite high ($> 85\%$) for rain, hail and snow, as particles are illuminated for only $25ms$. However, light throughput suffers from the long latency of the system making such system meaningless in strong snowstorms as it requires to switch off a majority of the rays. In mild or moderate rain/hail, the driver's visibility can still be improved significantly while losing maximum 20% of light.

Figure 11(b) shows the performance at $30km/hr$ of a more advanced system with the same parameters as the previous system, but with a $1ms$ exposure time and a $1.5ms$ total system latency. This can be achieved by using a gigabit interface and camera with a better quality sensor, and by eliminating the data transfer latencies by constructing an embedded and integrated imaging and projection system with a faster DLP projector. Notice the significant improvement in light throughput and accuracy in all weather conditions. This gives us hope that such a system may be realized in the near-future. Finally, it is interesting to compare these systems to the performance of a near-ideal system at $30km/hr$ that runs at $10kHz$ (Fig. 11(c)).

4. Prototype System for Reactive Illumination

As described by our simulator, a practical system for rain and snow will require high-speed imaging and projection with low latency for data throughput and processing. We expect that specialized embedded hardware with an integrated imaging and projection unit will be required to create a compact headlight. As proof-of-concept, we have developed a prototype system with off-the-shelf components to validate the findings of the simulations. The system consists of an optically co-located camera and projector operating at $120Hz$. A similar setup has been used for context-aware

lighting by Wang *et al.* [20], but their system is too slow for our application ($50-70ms$ latency). In order to provide repeatability and ground-truth, we use a precisely controllable artificial rain generator (Section 4.2). Processing was performed on $3.2GHz$ Intel Xeon processor with $8GB$ of RAM running Windows Vista 64-bit.

4.1. Camera and Projector Hardware

The experimental setup consisting of the co-located camera-projection prototype, drop generation hardware, and backdrop is shown in Figure 12. We use a monochrome camera (Point Grey Flea3) with ethernet interface that is capable of capturing $120FPS$ over a 120×244 region of interest. The projector is a Viewsonic PJD6251 DLP with a native resolution of 1024x768 and maximum vertical frequency of $120Hz$. It outputs a brightness of 3700 ANSI lumens, which is the equivalent of the popular D2S HID headlight (approx. 3200 lumens) [7].

The optical axes of the camera and projector are co-located by placing a 50/50 beam splitter in front of the camera and projector lenses as illustrated in Figure 12. The camera is mounted on a stage that permits fine-grained control over translation and rotation for calibration. Co-location of the camera and projector is achieved when shadows cast by the illuminated objects are no longer visible by the camera. Figure 12(c) shows an example image produced when the camera and projector are properly co-located.

4.2. Rain Generator

To test our reactive illumination system, we constructed a drop-generation testbed. We use a similar setup as [3, 8], where solenoid valves control drop formation. As shown in Figure 12, each valve is connected to a single emitter and receives slightly pressurized water from a tank elevated overhead. We use 16 direct-acting miniature solenoids (KIP 07420904), with a maximum flow rate of 0.03 cubic meters per minute. The drop emitters are physically located in the operating range of the system. Each emitter can release up to 60 drops per second. Though drop sizes will be different, the drop generation system can be controlled to generate the water equivalent of a range of precipitations — from a small drizzle to a large thunderstorm.

4.3. Detection and Prediction Algorithms

Our primary goal with respect to detection and prediction was to optimize our algorithms for execution time rather than accuracy. First, we perform background subtraction followed by blurring and thresholding to find the bright water spots. These spots are then segmented with connected-components. The drops from the rain-generator fall with near constant velocity at 16 pixels per frame, so prediction is straightforward. By carefully measuring our
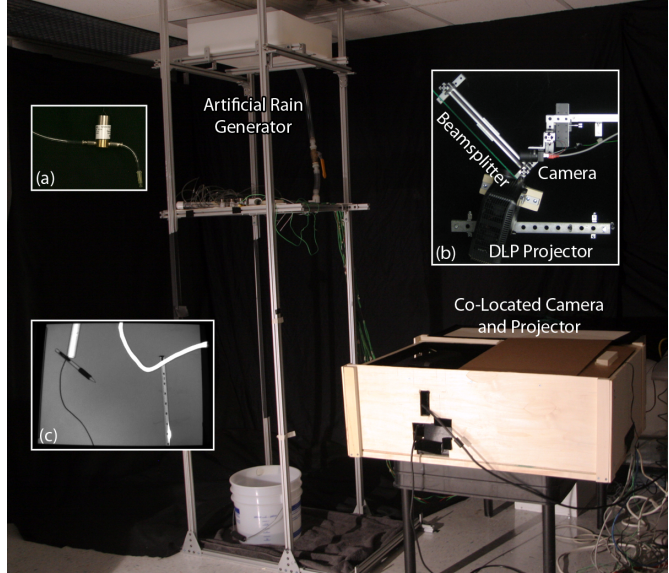

Figure 12. Our prototype and artificial rain generator. One of 16 direct-acting miniature solenoids for generating a water drop (a). The co-located system consists of a camera, projector, and beamsplitter (b). The camera is mounted on a stage that allows for fine adjustment. Example objects imaged by the co-located camera-projector pair. Notice the lack of visible shadows (c).

system latency, we can predict the image displacement of the drops over time.

Results of detection and prediction are shown for a single drop from an emitter in Figure 14 and for drops from all 16 emitters in Figure 15. Since the rain generator is completely controlled, experiments are repeatable and we can evaluate the system accurately by comparing the prediction locations to the actual locations of the drops made visible by illuminating them throughout the field of view. The single drop case shows how we can achieve high light throughput in a drizzle (99.7% throughput at 83.6% accuracy) and the 32 drops/s case demonstrates that our system is able to perform reasonably (98.1% throughput at 54.14% accuracy). Light throughput was calculated as the average for each frame.

The use of a beamsplitter reduces the light by 50% and reduces image contrast making it more difficult to detect and localize the water drops. The errors in prediction are mainly due to the errors in low-exposure drop detection at the top of the low-resolution frame and slight acceleration of the generated drops. The detection is the weakest link in our current system, but this can be improved with a higher quality sensor.

4.4. System Timing

The latencies of each stage in our system are illustrated in Figure 3. An exposure time of $5ms$ is used for the prototype. Other latencies were measured with a high-

resolution timer available in the Windows API and averaged (standard deviation) over 5500 frames. TX_c was $4.214ms$ ($0.182ms$), process was $4.081ms$ ($0.069ms$), and TX_p was $4.214ms$ ($0.182ms$). The time to transfer data to the projector ($TX_p = 9ms$) was measured by the following procedure. First, a DSP board (BF592 EZ-Lite) was used to illuminate an on-board LED and send a pulse to the DLP projector. Then, a high-speed camera (Photron 1024 PCI @ $500FPS$) was used to measure the time between observing the LED illuminate and the pulse project.

4.5. Performance and Limitations

The performance of our prototype with our rain test bed is shown in Figure 16. While the falling drops remain visible at the top of the picture (i.e. system latency), our prototype clearly improves visibility by successfully avoiding illumination of the drops once tracked. As we expected, not only are the drops less visible, the visual discomfort caused by falling raindrops is reduced and the overall visibility of the scene enhanced. Notice that even drops not illuminated remain *slightly* visible because of ambient lighting and drop detection inaccuracies. Light throughput was averaged over a 60 second duration. The mean light throughput is 97.4% with a standard deviation of 0.83%.

Despite the low resolution of 120×244 and the $13ms$ system latency, simulations using the same parameters as the prototype show that light throughput is still high when the platform is moving at $30km/hr$ (Figure 13). Note that compared to Figure 11(a), accuracy is lower but throughput is significantly higher because less particles are imaged for the settings used. Reaching such performance requires better detection and tracking.

Our current prototype requires two main improvements. First, in our experiments the imaging and illumination system is stationary while in actuality the vehicle would be moving. Second, our current prediction algorithm is not robust to chaotic wind. Finally, user studies in real rain and snow conditions are necessary to quantify the minimum accuracies and light throughput needed for a system to be successful based on human perception. Nonetheless, we believe that this work represents a strong initial step toward realizing a vehicle headlight that makes it easier to drive in bad weather at night.

5. Conclusions

In this paper, we analyzed the potential of using a high-speed co-located camera-projection system to reduce light scattering effects due to precipitation. Our simulation results show that it is possible to maintain light throughput well above 90% for various precipitation types and intensities. We then demonstrated a proof-of-concept prototype system operating at $120Hz$ on laboratory-generated rain
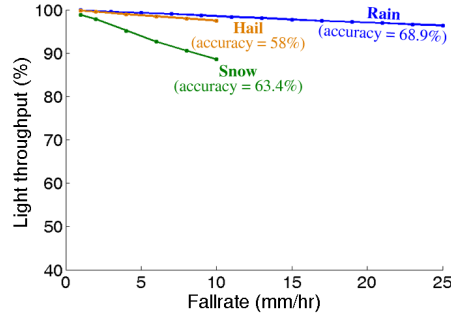

Figure 13. Performance the prototype system at $30km/hr$ could reach in rain, hail, and snow versus precipitation rate. Simulation shows that the light throughput remains high in all scenarios and that only 40% of the particles will be illuminated in the worst case, which should significantly improve driver's visibility.

that validates our initial simulations. Current LED lighting technology has already begun to adopt on-die sensing for color temperature correction. We believe that as this technology scales, it will soon be possible to realize an LED illumination source that is interleaved with an image sensor on a single chip. On-chip interconnects would virtually eliminate many of our critical timing latencies allowing for extremely high-speed operation at low-cost.

Acknowledgements

References

[1] D. Atlas, R. C. Srivastava, and R. S. Sekhon. Doppler radar characteristics of precipitation at vertical incidence. *Reviews of Geophysics*, 11(1):1–35, 1973. 4

[2] P. C. Barnum, T. Kanade, and S. Narasimhan. Spatio-temporal frequency analysis for removing rain and snow from videos. In *Proceedings of the First International Workshop on Photometric Analysis For Computer Vision - PACV 2007*, Rio de Janeiro, Brazil, 2007. 1

[3] P. C. Barnum, S. G. Narasimhan, and T. Kanade. A multi-layered display with water drops. volume 29, pages 76:1–76:7. ACM Transactions on Graphics, July 2010. 7

[4] N. E. Beam. Adaptive/anti-blinding headlights. Patent, 11 2000. US 6144158. 1

[5] A. S. Dennis, P. L. Smith, G. A. P. Peterson, and R. D. McNeil. Hailstone Size Distributions and Equivalent Radar Reflectivity Factors Computed from Hailstone Momentum Records. *Journal of Applied Meteorology*, 10(1):79–85, Feb. 1971. 4
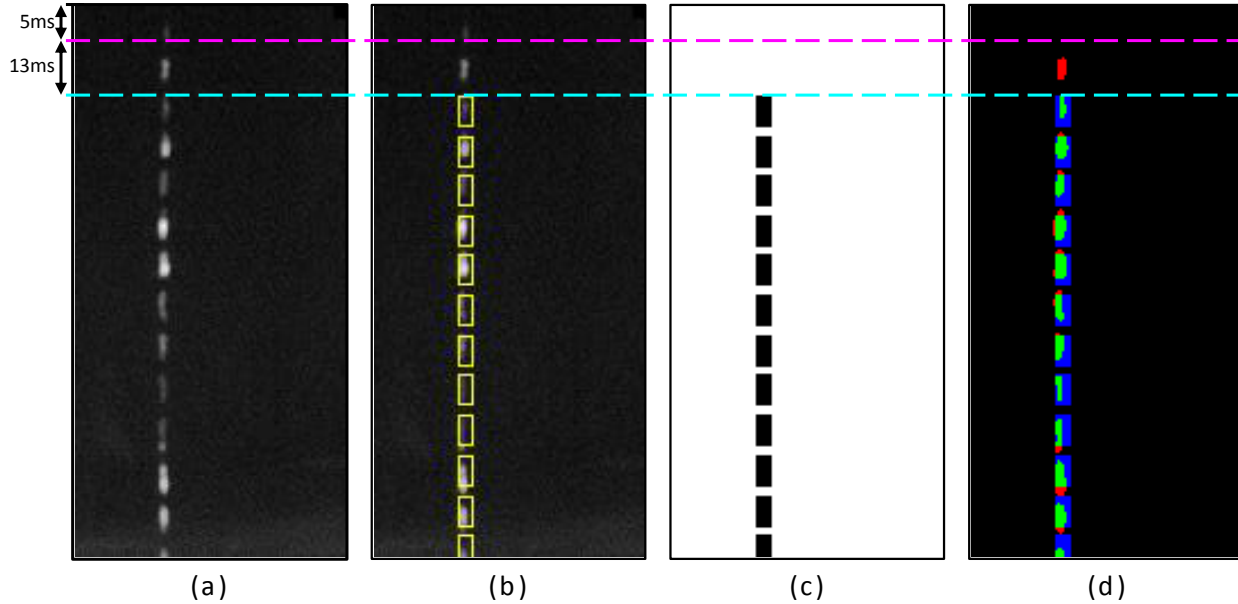
Figure 14. Predicting and avoiding illumination of a single drop. In this experiment, the drop is first imaged with $5ms$ camera exposure time (shown above magenta dashed line). The system latency is shown above the cyan dashed line. Images shown are composites of the 14 frames needed for the drop to traverse the FOV. Since the experiments are repeatable, we show the ground-truth image on the left with drops illuminated throughout the entire FOV (a). The drop is falling with near constant velocity at 16 pixels per frame, so prediction is straightforward (shown as yellow boxes) (b). The projected frames show the pattern the system uses to avoid illuminating the drop. (d) In this experiment, we achieved 99.7% light throughput and 83.6% accuracy. Red shows part of drop lit erroneously, green shows drop not illuminated correctly, and blue shows misclassified pixels (c).
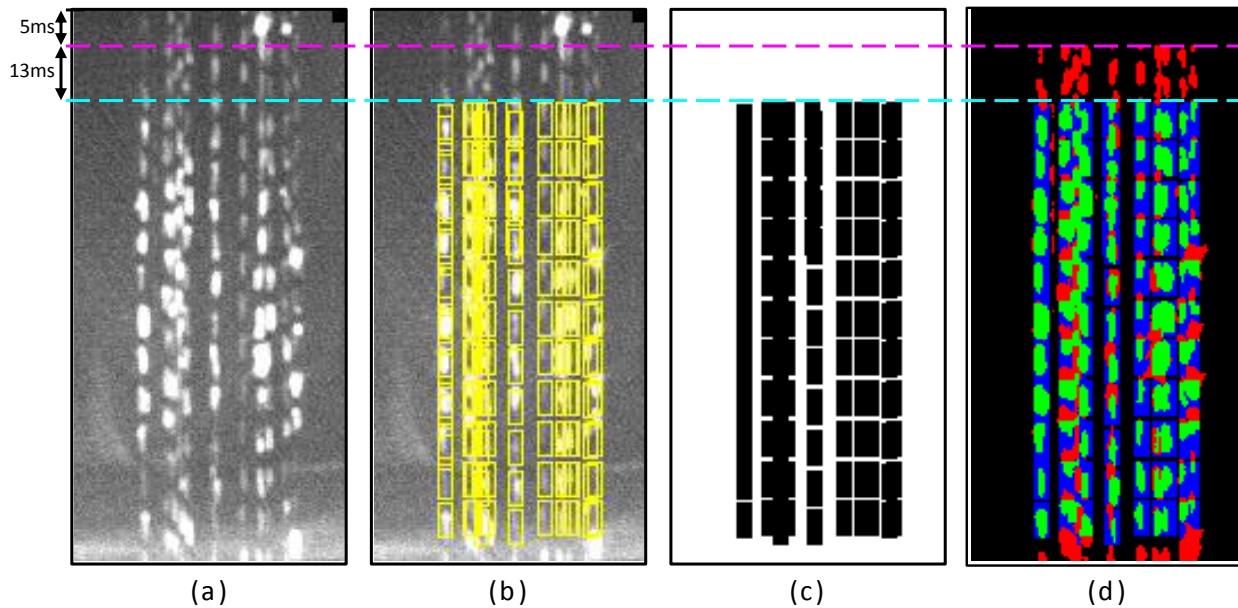

Figure 15. Predicting and avoiding illumination of 16 drops emitted at a rate of 32 drops per second. This is a difficult case with a water equivalent of a strong thunderstorm (over $90mm/hr$). Images shown are composites of the 65 frames needed for 16 drops to traverse the FOV. Our system is neither fast nor accurate enough to handle this scenario but it still performs reasonably (light throughput of 98.1%, and accuracy of 54.14%). The main bottleneck is the misclassification of pixels as drops, requiring more sophisticated detection and prediction.
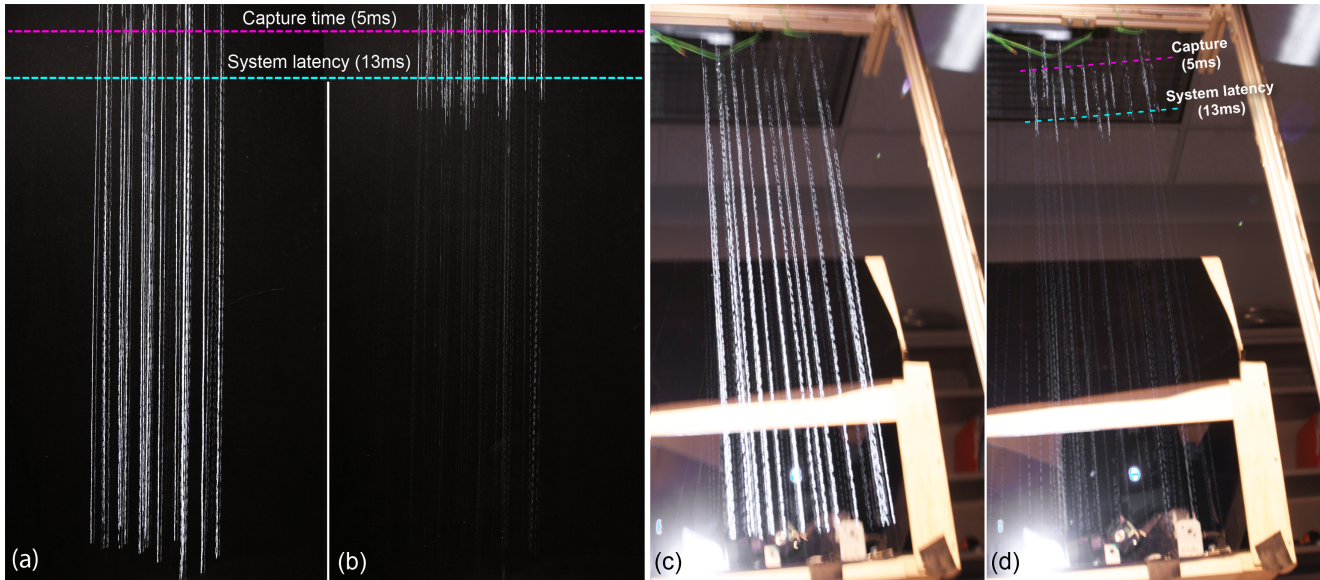
Figure 16. Naive illumination ((a) and (c)) versus fast reactive illumination ((b) and (d)) with our prototype running at $120Hz$ (system latency of $13ms$). Photos shown are captured with a $2.5s$ exposure time from two different viewpoints and backgrounds (a) and (b). The prototype successfully improves overall visibility by selectively turning off rays intersecting generated particles. The same repeatable rain is used in both experiments with 16 drops emitters generating drops at $2Hz$. Notice the impact of system latency (above cyan line) where drops are lit for a short period of time. Below this line our reactive illumination clearly reduces drops visibility and high frequency flickering patterns that cause visual discomfort. Even when tracked correctly, the drops remain *slightly* visible all along (notice the light streaks over dark background) due to prediction inaccuracies and ambient illumination.

[6] J. Dessens and R. Fraile. Hailstone size distributions in southwestern France. *Atmospheric Research*, 33(1-4):57–73, June 1994. 3, 4

[7] P. Dong. *Analysis, modelling and robust control of automotive HID headlight systems*. PhD thesis, The Hong Kong Polytechnic University, 2008. 7

[8] S. Eitoku, T. Tanikawa, and Y. Suzuki. Display Composed of Water Drops for Filling Space with Materialized Virtual Three-dimensional Objects. In *IEEE Virtual Reality Conference (VR 2006)*, pages 159–166. IEEE, 2006. 7

[9] K. Garg and S. K. Nayar. Detection and removal of rain from videos. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, volume 1, 2004. 1

[10] K. Garg and S. K. Nayar. When does a camera see rain? In *Proceedings of the Tenth IEEE International Conference on Computer Vision*, volume 2, 2005. 3

[11] K. Gunn and J. Marshall. The distribution with size of aggregate snowflakes. *Journal of Atmospheric Sciences*, 15:452–461, 1958. 3, 4

[12] N. Hautiere, E. Dumont, R. Bremond, and V. Ledoux. Review of the mechanisms of visibility reduction by rain and wet road. Technical report, 2009. 1

[13] J. A. Hull and J. J. Gallegos. Adaptive and interactive scene illumination. Patent, 06 2007. US 7227611. 1

[14] M. Langleben. The terminal velocity of snowflakes. *Quarterly Journal of the Royal Meteorological Society*, 80(344):174–181, 1954. 3, 4

[15] A. Macke and M. Groß klaus. Light scattering by nonspherical raindrops. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 60(3):355–363, Sept. 1998. 2

[16] J. Marshall and W. M. Palmer. The Distribution of Raindrops with Size. *Journal of the Atmospheric Sciences*, 5(4):165–166, Aug. 1948. 3, 4

[17] S. Nayar and S. Narasimhan. Vision in bad weather. In *The Proceedings of the Seventh IEEE International Conference on Computer Vision*, volume 2, 1999. 1

[18] Y. Shen, L. Ma, H. Liu, Y. Bao, and Z. Chen. Detecting and extracting natural snow from videos. *Information Processing Letters*, 110:1124–1130, Nov. 2010. 1

[19] J. H. Van Boxel. Numerical model for the fall speed of raindrops in a rainfall simulator. In D. Gabriels and W. Cornelis, editors, *Proceedings of the International Workshop on Technical aspects and use of wind tunnels for wind-erosion control, Combined effect of wind and water on erosion processes*, pages 77–85, Gent, 1997. 3, 4

[20] O. Wang, M. Fuchs, C. Fuchs, J. Davis, H.-P. Seidel, and H. P. A. Lensch. A context-aware light source. In *2010 IEEE International Conference on Computational Photography (ICCP)*, pages 1–8. IEEE, Mar. 2010. 7

[21] X. Zhang, H. Li, Y. Qi, W. K. Leow, and T. K. Ng. Rain Removal in Video by Combining Temporal and Chromatic Properties. In *Multimedia and Expo, 2006 IEEE International Conference on*, pages 461–464, Toronto, 2006. 1