Programmable Automotive Headlights

Srinivasa Narasimhan (PI)
Associate Professor, Robotics Institute, CMU

Contents

1 Research Team **2**

2 Report Summary **3**

3 Problem **4**

4 Approach **6**

5 Methodology **9**

5.1 Feasibility Study with Computer Simulations 9

5.2 Design Exploration 11

5.2.1 Dynamic Lighting 11

5.2.2 Improved Visibility in Snow 15

5.3 Programmable Headlight Prototype: Proof-of-Concept 17

5.4 Design and Implementation of a Prototype 19

5.4.1 Sensing the Road Environment 19

5.4.2 Image Processing and System Control 19

5.4.3 High-Speed Illumination of the Road Environment 21

5.5 System Calibration 21

5.6 Measuring System Latency 23

5.7 Programmable Headlight Prototype: Faster and More Compact . . 24

5.7.1 Imaging 24

5.7.2 Projection 26

5.8 System Response Time 27

5.8.1 Measuring End-to-End Latency 27

5.8.2 Effect of Prediction Algorithms on Latency 27

6 Findings **33**
6.1 Glare Free High Beams 33
6.1.1 System Requirements and Comparisons 33
6.1.2 Our Headlight Design as Glare Free High Beams 34
6.1.3 Road Demonstrations 35
6.2 Demonstrating Headlight Versatility 37
6.2.1 Improving Visibility During Snowstorms 38
6.2.2 Improved Lane Illumination 39
6.2.3 Early Visual Warning of Obstacles 41

7 Conclusions **43**

8 Recommendations **44**

Chapter 1

Research Team

Current:

- Andrew Rowe, Associate Professor Electrical and Computer Engineering, CMU
- James Hoe, Professor Electrical and Computer Engineering, CMU
- Robert Tamburo Project Scientist, CMU
- Srihari Sankar MS, Electrical and Computer Engineering, CMU
- Marie Nguyen PhD, Electrical and Computer Engineering, CMU

Past:

- Takeo Kanade Professor, Robotics Institute, CMU
- Mei Chen Intel Labs
- Eriko Nurvitadhi Intel Labs
- Subhagato Dutta MS, Electrical and Computer Engineering, CMU
- Zisimos Economou Research Staff, CMU
- Abhishek Chugh MS, Electrical and Computer Engineering, CMU
- Prakaash Karthikeyan MS, Electrical and Computer Engineering, CMU
- Vivek Umapathi MS, Electrical and Computer Engineering, CMU

Chapter 2

Report Summary

The primary goal of an automotive headlight is to improve safety in low light and poor weather conditions. But, despite decades of innovation on light sources, more than half of accidents occur at night even with less traffic on the road. Recent developments in adaptive lighting have addressed some limitations of standard headlights, however, they have limited flexibility - switching between high and low beams, turning off beams toward the opposing lane, or rotating the beam as the vehicle turns - and are not designed for all driving environments. We introduce an ultra-low latency headlight that can sense, react, and adapt quickly to any environment while moving at highway speeds. Our single hardware design can be programmed to perform a variety of tasks. Glare free high beams, improved driver visibility during snowstorms, increased contrast of lanes, markings, and sidewalks, and early visual warning of obstacles are demonstrated.

This report is divided into five main parts. In the first part (Chapter 3), the problem addressed is described. In Chapter 4, our approach to the problem is described. In Chapter 5, we use computer simulations to examine the feasibility (latency requirements) of our headlight design. Lessons learned from the design exploration were used to guide the development of headlight prototypes (Sections 5.3 and 5.7). In Chapter 6, a variety of automotive applications are demonstrated on the road with the prototype. Finally, in Chapters 7 and 8, conclusions and recommendations are provided.

Chapter 3

Problem

Traditional headlights consist of a small number of lamps with simple optics to direct a light beam onto the road. Starting with gas/oil lamps in the 1880s, research has been primarily geared towards developing headlights that can be electrically controlled, have a long working life, and are bright and energy efficient. The inventions of Halogen lamps, Xenon (HID) lamps [9], [20], and the more recent LED [19], [22] and Laser sources [26] have followed this research trend. These latest sources provide bright and comfortable color temperatures improving driving experiences. However, even with these new light sources the only control offered to a majority of drivers is to switch between high and low beams.

Low beams illuminate the road a short range in front of the vehicle while high beams have a longer range and wider angle. High beams are useful in a variety of situations providing better visibility farther down the road and along narrow, curvy roads. However, they cause significant glare to other drivers, bicyclists, and pedestrians. Glare is especially troublesome for older drivers causing a 55-year old 8 times as long to recover than an 18-year old [13]. High beams also significantly reduce contrast in the presence of fog and haze, and cause bright distracting streaks during precipitation events. In 2013, there were more than 300,000 crashes and thousands of fatalities caused by rain and snow [7].

Even after 130 years of headlight development, more than half of vehicle crashes and fatalities occur at night despite significantly less traffic [7]. Approximately 30% of drivers are stressed by glare causing many fatalities every year [5]. Moreover, more there are more than 270,000 crashes with animals, 31,000 crashes with pedestrians, close to 15,000 crashes with bicyclists, and thousands of fatalities from veering off the lane and road [7] (Fig. 3.1).

Recognizing the limitations of traditional headlights, adaptive lighting sys-

TOTAL NUMBER AND PERCENTAGE OF FATALITIES AND CRASHES AT NIGHT INVOLVING VEHICLES

	Bicyclists	Pedestrians	Rain	Snow/Sleet	Animals*	Running off Road or Lane*	Obscured Vision*
Fatalities	360 (49%)	3,300 (70%)	1,200 (56%)	340 (51%)	180	3,720	1,490
Crashes	14,360 (29%)	31,300 (44%)	208,000 (36%)	111,000 (45%)	273,000	NA	NA

*Includes daytime statistics

National Highway Traffic Safety Administration, "TRAFFIC SAFETY FACTS 2013"
AAA Foundation for Safety Research, "How to Avoid Headlight Glare"
National Highway Traffic Safety Administration, "DRIVERS' PERCEPTIONS OF HEADLIGHT..." 2003.

Figure 3.1: Traffic safety statistics during night time driving.

tems have been developed to adjust their brightness in response to changing driving conditions. Some systems, e.g. Lincoln [10], Audi [22], Volkswagen [17], mechanically swivel the headlight based on the vehicle's turning radius allowing drivers to see around curved roads. Other systems use configurations of multiple LEDs, where individual LEDs can be automatically turned off toward the driving lane and/or the opposing lane to reduce glare, e.g., BMW [27], Audi [18], Mercedes [16], and Volvo [14]. In [26] swiveling LEDs spotlight pedestrians on sidewalks. These advanced systems have come a long way from traditional headlights, but fundamental issues remain: they are not versatile and are designed for one-off applications, they require mechanical problems that reduce reliability, and their low-resolution and high latency limits them from adapting to many types of road conditions and poor visibility situations. Thus, a headlight that adapts to the environment can be critical to improving safety on the road during poor visibility conditions.

Chapter 4

Approach

We present a new design for an automotive headlight that is flexible and can be programmed to perform multiple tasks at high speeds. As shown in Figure 4.1, the key idea is the introduction of a high-resolution spatial light modulator (SLM) such as the digital micro-mirror device (DMD) present in DLP projectors. A DMD divides a light beam into approximately one million that can be individually controlled to shape the collective beam for any situation. A sensor (camera) is co-located with the light source and a computer processes images to generate illumination patterns for the SLM. While the design may seem straightforward, there are many challenges in building such a system to serve as a headlight.

The accuracy requirements can be high since small errors in beam positioning and flickering are easily perceived and can be more disturbing than standard headlights. High accuracy can be achieved by minimizing the time from when a camera senses the environment to when the headlight reacts (system latency). Low latency is also required to avoid the need for complex prediction algorithms to determine where an object will move next. For example, assume that the scene (physical environment) consists of multiple objects that are moving independently at differing speeds and directions. The goal of our headlight would be to sense (image) the scene, detect objects and predict their future locations, and illuminate or dis-illuminate the objects (Fig. 4.2). Predicting the location of randomly moving objects with a high latency system is challenging. However, if the system's latency is very low (i.e., reaction time is very short), the complexity of the prediction algorithm can be much simpler. In the following section, we investigate the feasibility of our design, with computer simulations, by characterizing its real-time performance in terms of latency for given prediction algorithms that have an inherent tradeoff between execution time and accuracy.

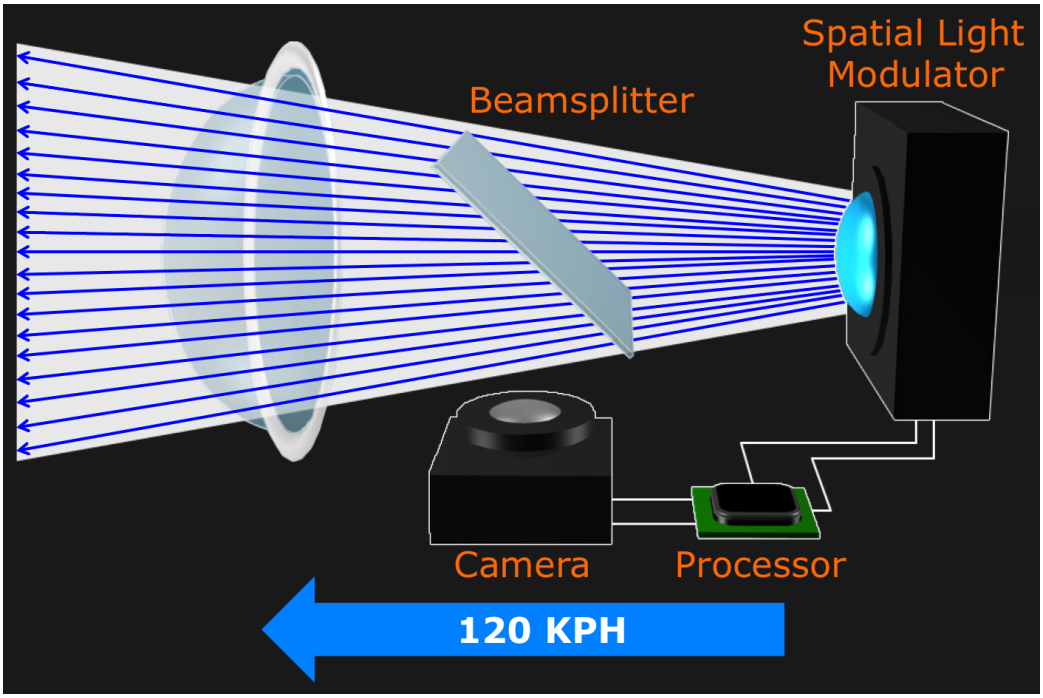

Figure 4.1: Our versatile headlight design utilizes a high-resolution spatial light modulator permitting the fine control of light rays in reaction to any number and type of objects detected from captured images.

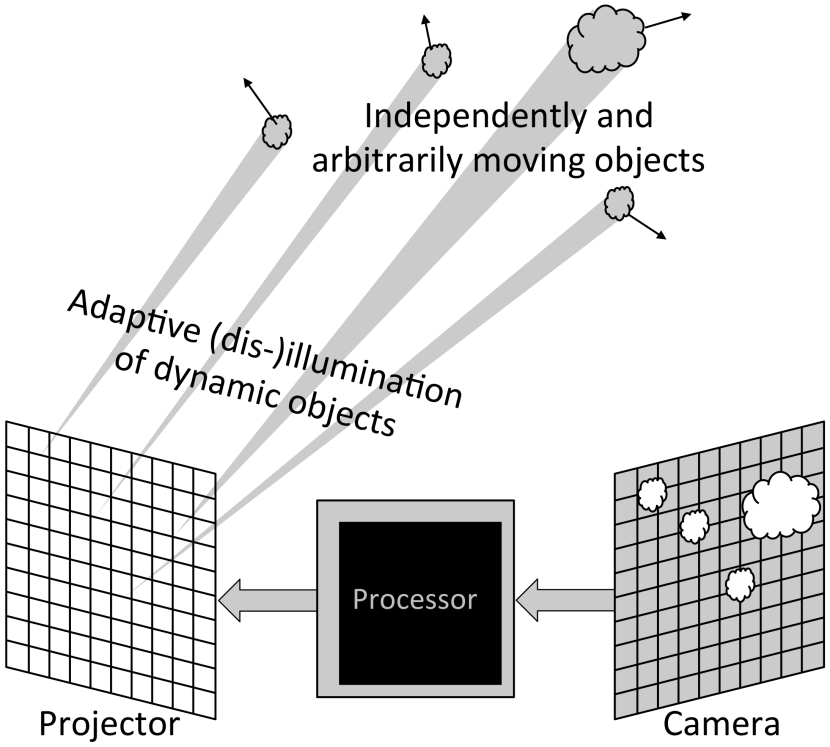

Figure 4.2: The headlight illuminates or dis-illuminates objects detected from images captured by a camera.

Chapter 5

Methodology

5.1 Feasibility Study with Computer Simulations

We present the evaluation of our low-latency, high-speed programmable headlight under various application workloads. We characterize the system's Quality-of-Service in terms of *error* and *light throughput* as a function of the system's response time. We have developed a simulator that allows us to explore the design space of such a system, including various strategies for predicting object movement. We identify key operating points where decreasing the system latency with a less complex algorithm can outperform higher complexity algorithms that take more time to run to achieve the same goal. We see a similar performance trade-off where jitter (not just delay) in the control loop decreases performance.

The configuration of a Programmable Headlight is shown in Fig. 5.1. The camera captures an image of a scene containing a set of objects with unknown locations and arbitrary motion. The processor then identifies the objects and records their current locations. This process is termed *binarization*, where, image pixels corresponding to the object locations are 'on' and are 'off' otherwise. The binary image, and perhaps a history of previous binary images, is then used to predict the object locations at the end of the system's response time. The predicted image is then warped to the coordinate frame of the projector, which in turn illuminates or dis-illuminates (depending on the application) the detected objects by turning pixels on or off, respectively. Note that while we focus on binary systems, pulse-width modulation coding can be used to project arbitrary grayscale values.

In order to compute the mapping between light rays from the objects observed by the camera and the corresponding light rays exiting the projector towards the

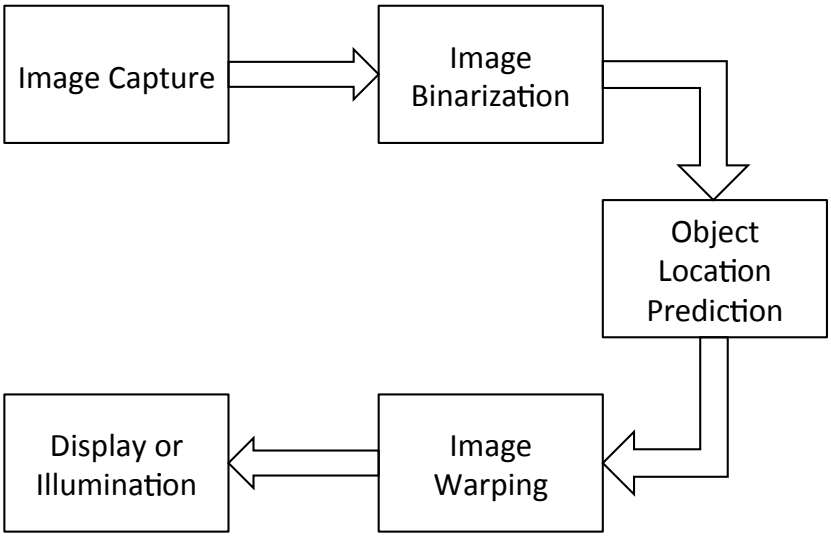

Figure 5.1: Block diagram of main architectural components of a Programmable Headlight.

objects, the three-dimensional positions of the objects must be computed. Such a mapping is commonly achieved by using multiple cameras in a stereoscopic configuration. Instead, for the purposes of this work, the camera and projector are optically co-located using a beamsplitter, which avoids parallax between the camera image and the projector image [23]. Co-location removes the need to compute the distances of objects in three dimensions and allows for all computations to be performed in the image space.

The latency (response time) of a Programmable Headlight is defined as the time elapsed from the start of exposure by the camera and the completed illumination by the projector. In addition to measuring latency, we also measure the jitter or uncertainty in the end-to-end latency. In this work, systems (simulated and hardware prototype) with varying parameters are evaluated based on the following performance measures as a function of the system's latency and jitter. *Error* quantifies the incorrect illumination or dis-illumination of the scene and *light throughput* quantifies the total amount of light that illuminates the scene.

We apply this evaluation to two types of applications. The objective of the first application is to dynamically illuminate objects while the objective of the second application is to dis-illuminate objects. For the illumination application, we use rigid objects moving in a linear or projectile motion (ping pong balls moving and colliding). System error for this application is computed as the percentage of

pixels that do not illuminate the objects and illuminate the background. For the dis-illumination application, we consider a large number of small objects moving chaotically, such as snowflakes during a storm. This application evaluates how well the system can act to avoid illuminating snowflakes to reduce their visibility to observers or a camera. In this case, error is computed as the percentage of pixels that incorrectly illuminate the objects. The error caused when incorrectly dis-illuminating the scene (road environment) is calculated separately as light throughput to assess the trade-off between reducing snowflake visibility and illuminating the road for the observer.

5.2 Design Exploration

Our simulator models multiple types of objects and their motions in the scene to characterize the performance of programmable headlight systems with different latencies performing different algorithms. The scene is rendered using OpenGL and the rendered image is an input to the simulated system, which analyzes the image and generates a corresponding illumination pattern. For the remainder of this section, we explore, for the first time, the performance of programmable headlight systems for an illumination and a dis-illumination application.

5.2.1 Dynamic Lighting

Dynamic illumination of moving objects could be used, for example, to spotlight pedestrians, bicyclists, construction workers, wildlife, etc. The actual system illuminates the scene with infrared light sources and observes the environment with a near-infrared, monochrome camera ensuring that the system's output (visible light) is not captured by the camera. The input image is thresholded to produce a binary image, which can then be displayed for immediate system response. We refer to this as the *no prediction algorithm*. Alternatively, prediction strategies can be employed to compensate for object motion - most likely adding to the system's latency though.

A simple algorithm for motion compensation assumes objects have an equal probability of moving in any direction. This algorithm is implemented by performing dilation with a kernel of fixed size on the binary image. A more intelligent approach would predict the future location of individual objects and produce the illumination pattern based on this information. In our implementation, blobs are detected from binary images and stored for the two most recent images. A

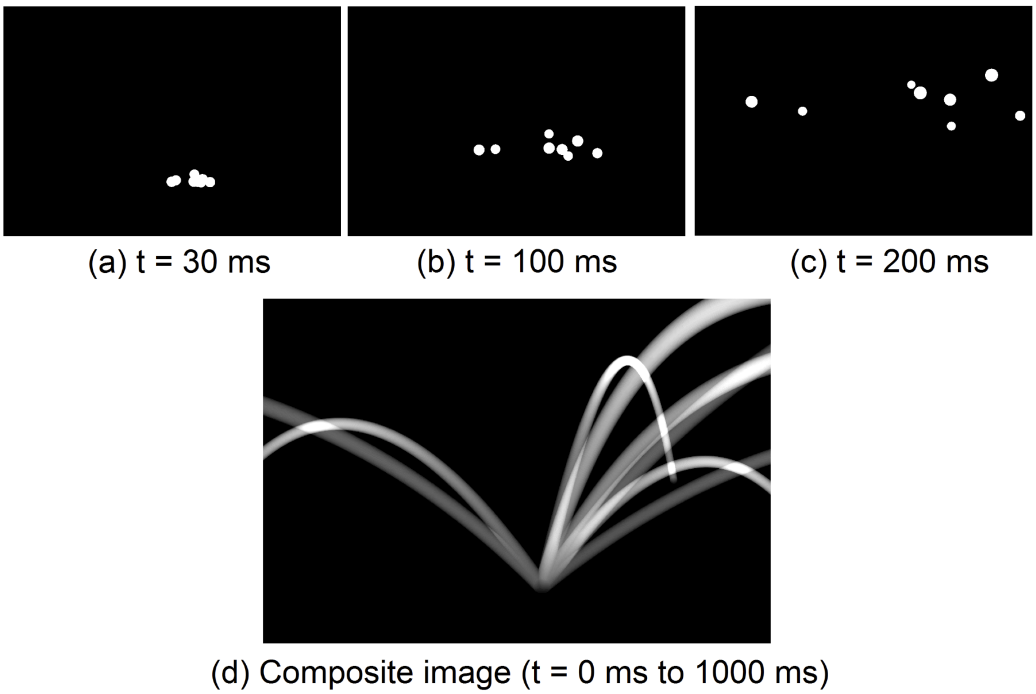

Figure 5.2: Images (a)-(c) are renderings of rigid objects (e.g., ping pong balls) quickly accelerating from rest (exploding) at different times. A long exposure (1 s) image of the event is shown in (d).

metric based on blob size and position is used to find correspondences between newly detected blobs and previously detected blobs. The image velocity and position of objects in the two images are used to linearly extrapolate the position of objects at the time of illumination to produce an illumination pattern.

Through simulation, system performance was evaluated for these different prediction strategies. For performance evaluation, the fast motion of multiple rigid objects (perhaps, ping pong balls) being struck with a larger object (perhaps, a tennis ball) was simulated (Fig. 5.2). After receiving an image of the event, the simulated programmable headlight system generates a binary illumination pattern. Based on the system's latency, an image of the scene is generated the instant that it is dynamically illuminated.

For quantitative evaluation, the image of the scene at the time of illumination and the illumination pattern are compared. There are two types of error; *Positive Error* and *Negative Error*. Positive Error is caused when a foreground object is not

correctly illuminated and can be calculated as Positive Error $= \frac{\text{Object pixels not illuminated}}{\text{Total object pixels}}$. Negative Error is caused when the background is incorrectly illuminated and is defined as Negative Error $= \frac{\text{Background pixels illuminated}}{\text{Total object pixels}}$. The Performance of the system for this application is therefore measurable by Total Error $= \frac{\text{Positive Error + Negative Error}}{2}$.

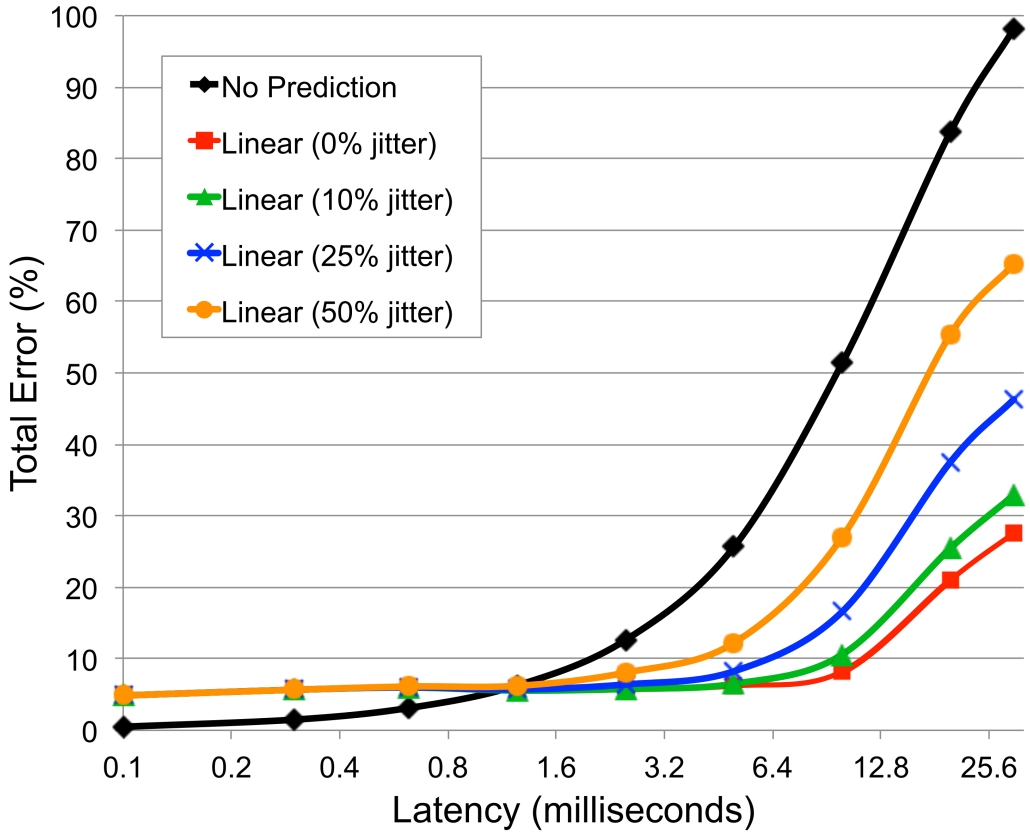

Figure 5.3: Total error while illuminating rigid objects as they accelerate for different system latencies. The black curve shows errors achieved with the no prediction algorithm. The other curves show error with linear prediction and latency uncertainty. Latency is plotted on a logarithmic scale.

As shown in Fig. 5.3, the no prediction algorithm has an error approaching 100% for latencies of typical video frame rate (30 Hz). Error can be reduced dramatically to around 10% by decreasing the system latency to 2 ms or less. At this latency, the no prediction algorithm is comparable to the linear prediction algorithm, making the no prediction algorithm a viable strategy for fast reacting systems. For slower systems, the linear prediction algorithm results in much bet-

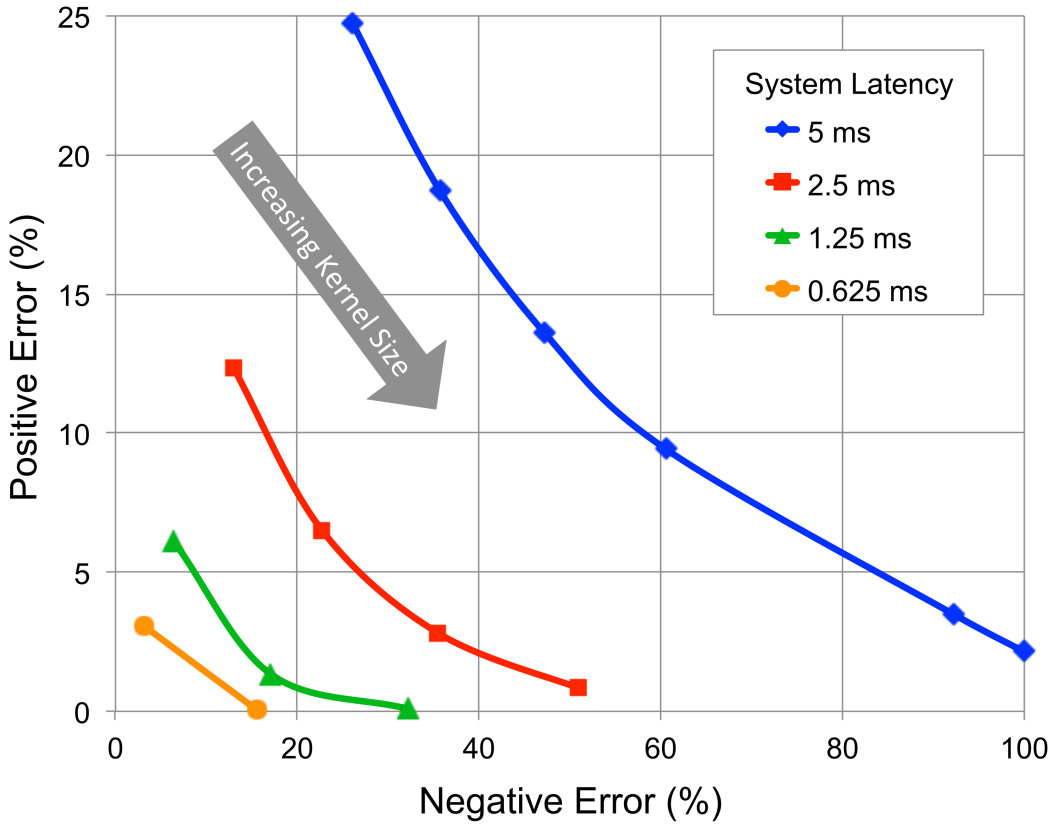

Figure 5.4: Increasing dilation kernel size improves proper illumination of the objects (positive error decreases) while increasing incorrect illumination of the background (negative error increases).

ter performance. The downside of the linear prediction algorithm are sources of error that arise from incorrect correspondences, especially when the objects are near each other or overlap. Even at low latencies, this type of error is difficult to remove. At high image capture rates (e.g., 1 kHz exposure time), it is better to avoid linear prediction since typical objects will not move much in such short time periods. The contrary may be true for objects that move extremely fast. Since the prediction algorithm uses system latency to predict the position of objects at the time of illumination, it is imperative to study the effect of jitter. In the simulator, jitter is modeled as Gaussian noise in the latency. In Fig. 5.3, we also compare systems with different amounts of jitter (standard deviation in Gaussian distribution) and, although, jitter is an important consideration at higher latencies,

its effect is minor at lower latencies.

The effect of using dilation with kernels of varying size to compensate for object motion is shown in Fig. 5.4. By increasing the size of the kernel, objects are better illuminated (decreasing positive error), but with the tradeoff that more of the background is incorrectly illuminated (increasing negative error). Compared to the no prediction and linear prediction algorithms, dilation results in more total error even with the smallest kernel size. Our simulations show that, for dynamic lighting of relatively fast moving objects, there is no advantage to using a prediction algorithm for low latency systems. Whereas, higher latency systems greatly benefit with linear prediction.

5.2.2 Improved Visibility in Snow

Retro-reflection from falling snowflakes distracts drivers from observing the road and makes driving during a snowstorm dangerous and stressful. Using our headlight, snowflakes can be dis-illuminated for a very short period to improve visibility. We simulate snow falling at 3 mm/hr [12] and study the performance of different systems (Fig. 5.5). Note that the objects (snowflakes) in this application are small in size but large in number. The objective with this application is opposite of the dynamic lighting application, i.e., objects are dis-illuminated instead of illuminated. The two performance metrics to characterize system behavior for this application are Error $= \frac{\text{Snow pixels illuminated}}{\text{Total snow pixels}}$ and Light Throughput $= 1 - \frac{\text{Pixels dis-illuminated}}{\text{Total pixels}}$.

The average light throughput with each algorithm was 95.7% for no prediction, 84.9% for linear prediction, and 87.6% for dilation with a 1×5 kernel. As shown in Fig. 5.6, the system performs well with the no prediction algorithm at lower latencies. Performance degrades for systems with 0.2 ms latency as compared to both the dilation and linear prediction algorithms. For systems with a latency below 2 ms, dilation with a vertical kernel of size 1×5 performs much better than linear prediction while maintaining higher light throughput. Linear prediction performs best for systems with a latency greater than 2 ms.

Although, it is clear that dilation reduces error, the effect of dilation on light throughput is investigated. Simulations were performed to measure light throughput for achieving 0% error. Results show that systems with latency below 1.6 ms have more than 90% light throughput, but systems with longer latency have less than 85% light throughput. Additionally, the amount of dilation is also dependent on the speed of the snowflakes. For example, if we wish to dis-illuminate snowflakes in a strong wind or on a fast moving vehicle, dilation would become inappropriate with even low latency systems.

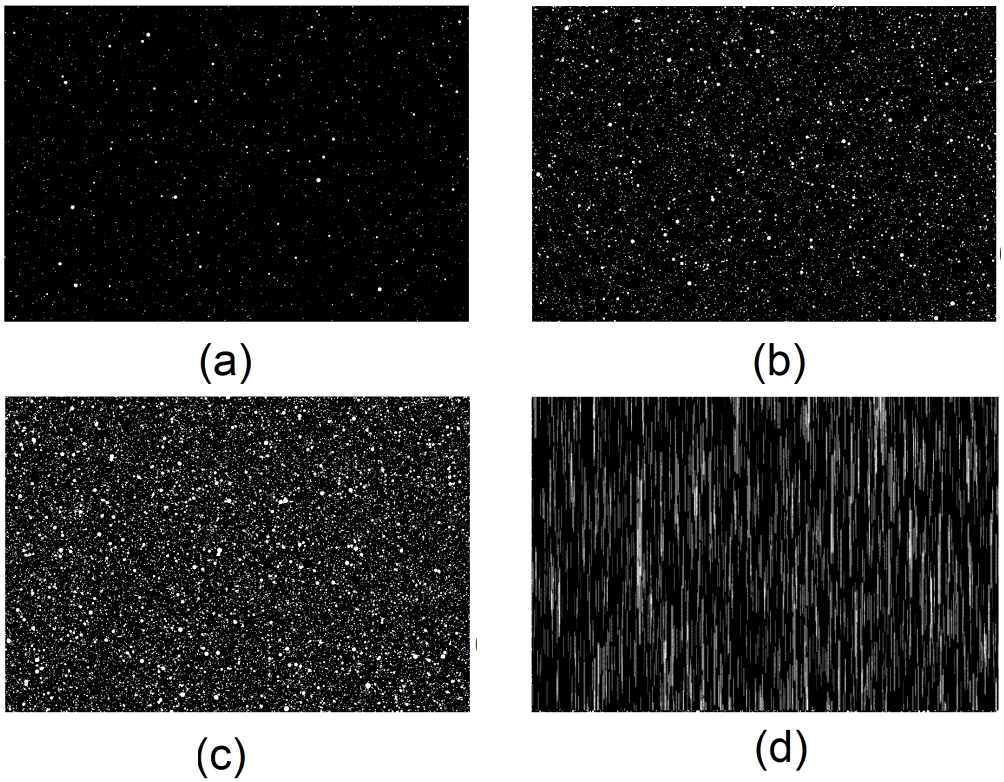

Figure 5.5: Images rendered by the simulator for snowflakes falling at (a) 0.5 mm/hr (b) 3mm/hr (c) 10 mm/hr. (d) is a long exposure (60 ms) capture of snow falling at 0.5 mm/hr.

We performed simulations to understand the effect of snowfall rate for our implemented prototype described in Sections 5.7 and 5.3. Latencies measured from our system were used to compare system error with different algorithms. Changes in snowfall rate do not change the latency of the systems with the no prediction or dilation algorithm. Thus, their error rates do not change either with the no prediction algorithm resulting in 70% error and the dilation algorithm resulting in 2% error. On the other hand, the linear prediction algorithm depends on the number of objects in the image, and therefore, causes an increase in computation time resulting in more error for higher snowfall rates. Error is approximately 15% for a light snowfall (0.3 mm/hr) and increases to 35% for blizzard-like snowfall rates (10 mm/hr).

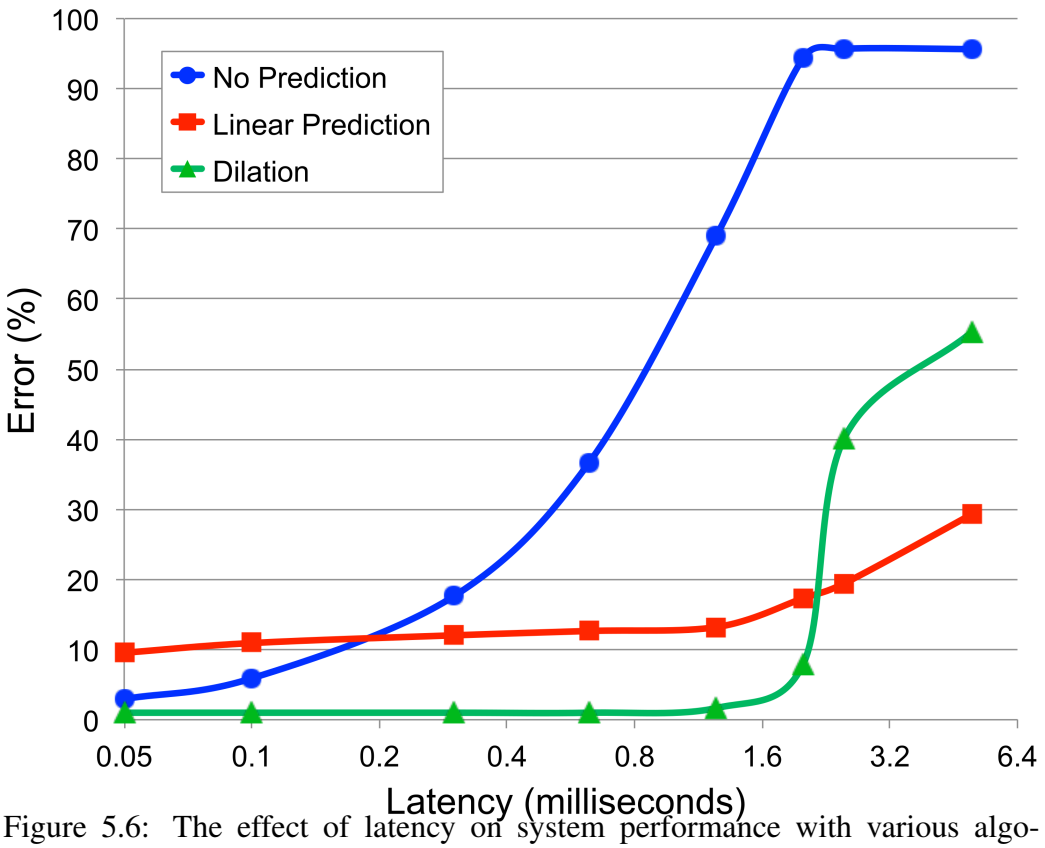

Figure 5.6: The effect of latency on system performance with various algorithms. On average, the no prediction algorithm achieves 95.7% light throughput, while linear prediction and dilation (1×5 kernel) achieve 84.9% and 87.6% light throughput, respectively. Latency is plotted on a logarithmic scale.

5.3 Programmable Headlight Prototype: Proof-of-Concept

The lessons learned from design exploration in the previous Section were used to guide the development of a Programmable Headlight prototype, which consists of four main components: an image sensor, processing unit, spatial light modulator (SLM), and beam splitter. The *imaging sensor* observes the road environment in front of the vehicle. Additional sensors such as RADAR or LIDAR can be incorporated into the design to complement the camera. The *processor* analyzes image data from the sensor and controls the headlight beam via a spatial light modulator. The *spatial light modulator* (e.g., digital micro-mirror device, liquid crystal display, liquid crystal on silicon, etc.) modifies the beam from a light

source by varying the intensity over space and time in two dimensions. We use a DMD because its high working frequency and small pixel size permit high-speed modulation and fine illumination control, which makes it possible for our headlight to quickly react to objects as small as snowflakes and objects as large as vehicles.

The camera and SLM are co-located along the same optical line of sight via a *beam splitter*, which virtually places the image sensor and DMD at the same location. Co-location is advantageous because it makes calculating the distance to objects unnecessary. Consequently, there is no need to perform costly computations required for depth estimation and 3D tracking. Also, a single homography will map the camera and projector image planes regardless of the scene. If the image sensor and DMD chip are placed very close to each other, the beam splitter is not required. Similar systems with a similar design have been described by [25] and [24], but their systems are too slow for high-speed automotive applications. High latency in conjunction with road effects like wind turbulence and vibration will require complex prediction algorithms that will add latency to the system making it unusable.

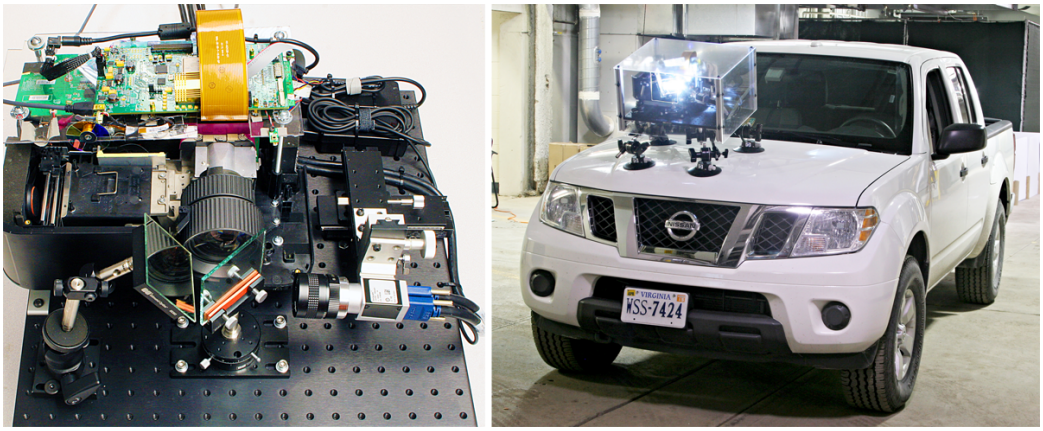

Figure 5.7: Left: Prototype of our programmable automotive headlight design (computer not pictured). The camera, spatial light modulator, and beam splitter are firmly mounted to an optical breadboard. A mirror to the side of the beam splitter deflects reflected light from the light source upward. Right: Road tests were conducted by securing the prototype to the hood of a vehicle with a suction cup-based mount. An acrylic enclosure was constructed to protect components from dust, dirt, and moisture.

5.4 Design and Implementation of a Prototype

We designed and implemented a prototype with low latency and high data throughput (Figure 5.7), and conducted road tests to demonstrate the feasibility of our DMD-based headlight design. The camera and SLM must have a very fast frame rate, e.g., kilohertz range, to capture images of fast moving objects and to create illumination patterns that are imperceptible to drivers. Consequently, a lot of data must be transferred to and from the processing unit with minimal latency. To achieve these goals, components with high-speed interfaces were tightly integrated through hardware and software. The prototype measures 45 cm wide, 45 cm long, and 30 cm tall and is currently too large to install in a vehicle as a headlight. The current size is due to using off-the-shelf components. Specialized embedded hardware with an integrated imaging, processing, and SLM unit will be required to create a compact headlight. Road tests were conducted by securing the prototype to the hood of a vehicle with a suction-cup based vehicle mount. A custom acrylic enclosure protects the system from dust, dirt, and moisture. We demonstrate in Sections 6.1 and 6.2 that the prototype performs a variety of tasks at typical traffic speeds.

5.4.1 Sensing the Road Environment

A camera (Basler acA2040) with a CMOS sensor highly sensitive to light with correlated double sampling to significantly reduce noise was used to capture images. The camera is sensitive to visible and near infrared light since most objects of interest are detectable within this spectrum of light. Monochrome imagery is used to avoid the computational overhead associated with demosaicing the Bayer pattern. A global shutter with area scan is used to avoid distortion effects common with the rolling shutter. Latency is reduced via a pipelined pixel architecture that permits exposure during readout. The camera's extended CameraLink configuration has transfer rates of up to 6.8 gigabits per second. The camera is mounted to a set of linear stages for fine control during calibration.

5.4.2 Image Processing and System Control

A desktop computer provides an interface between the camera and SLM, performs image analysis, and controls the system. The computer was custom built using an Intel Core 3.4 GHz (i7-2600K) CPU with eight cores and hyper-threading technology. A PCI express 2.0 frame grabber (Bitflow Karbon SP) that transfers image

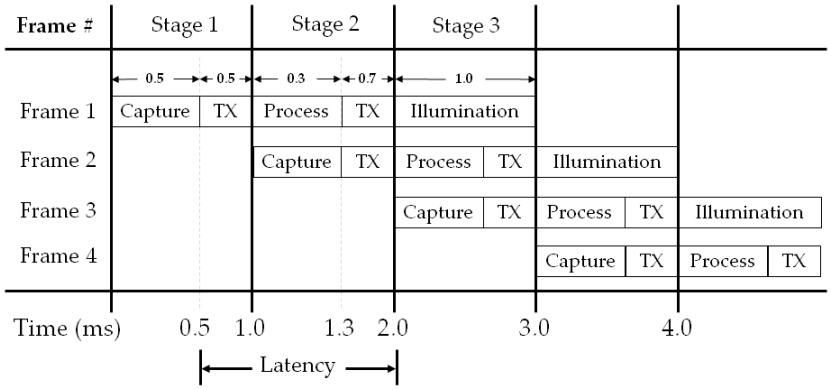

Figure 5.8: Timing diagram of the three-stage pipeline with execution times in milliseconds. Capture refers to camera exposure. Process refers to analysis of images and system control. TX denotes data transfer between camera and computer or between computer and SLM. Latency is the time required to illuminate the scene after capturing an image.

data directly into computer memory without any buffering. The main processing tasks were parallelized to reduce latency and increase system responsiveness. The three-stage processing pipeline is shown in a timing diagram (Figure 5.8) with times measured from the prototype system as described in Section 5.6. Capture refers to the integration time of the camera. TX denotes the time to transfer image data to the host computer and the time to transfer data from the host computer to the SLM. Process refers to image analysis and system control. Illumination refers to directing light to the scene for a single cycle.

Since execution time is critical, the focus of image analysis algorithms is on speed rather than accuracy. Image analyses were performed using OpenCV compiled with Intel Integrated Performance Primitives and Thread Building Blocks to maximize parallelism. Functions that perform per-pixel operations were combined using SSE2 intrinsic functions, when possible, to reduce the computation time associated with multiple iterations over the image. Pre-computable operations such as distortion correction and perspective transformation were initialized and stored in look-up tables. After analyzing images, illumination patterns are encoded and stored in an array then transmitted to the SLM.

5.4.3 High-Speed Illumination of the Road Environment

A DMD chip is used as a SLM for its spatial and temporal resolution. They are used in consumer DLP projectors, but are driven by video frame rates, which are well below our kilohertz target. A DLP development kit (WinTech W4100) based on the Discovery 4100 (Texas Instruments) was used as the basis of our SLM because the board contains a user programmable FPGA (Xilinx Virtex-5) to achieve fast update rates. The DMD chip is 0.7" with XGA (1024×768) resolution, which, essentially means the headlight beam can be divided into 786,432 smaller beams each of which can be turned on or off. This type of modulation gives unprecedented control over the illumination in space and time. Illumination patterns are received from the host computer by USB 2.0.

The DLP development kit does not include any optics. The optics and light source from a consumer DLP projector (InFocus IN3124) were used instead of designing custom components. We chose this projector because it uses the same DMD chipset as the development kit and uses a lamp brighter (4800 Lumens) than most vehicle high beams. The projector's native DMD chip was removed from the optics module and replaced with that of the development kit via a custom machined mount. A copper heat sink and fan were installed to improve heat dissipation. All of the native DLP electronic boards were left attached to maintain operability even though only the optics and lamp are actively used.

The FPGA was programmed to display patterns faster than 1 kHz. In our design, the FPGA receives data streamed from the host PC and produces the commands for a DMD controller to display the appropriate patterns on the DMD. Each row (1024 pixels) of the DMD is represented as a bit-vector. Transferring a 1024-bit vector for each of the 768 rows was too slow (over 1.5 ms). Instead, the rows are subsampled by a factor of four by representing 1024 pixels by a 256-bit vector. Some resolution is lost, but the visual impact is negligible. Data was further compressed to increase system speed by reading out every other row from the image sensor. The missing rows of the resulting illumination pattern are filled-in by duplicating the previous row on the FPGA. Thus, the image is down-sampled by a factor of 4 horizontally and a factor of 2 vertically.

5.5 System Calibration

Calibrating the system consists of co-locating the camera and SLM, and computing the homography between the camera and SLM image planes. To achieve

co-location, a beam splitter with 50% transmission and 50% reflection (Edmund Optics) is used. The projector, rigidly affixed to the optical breadboard, illuminates an object. The camera is translated in all three cardinal directions and rotated until shadows cast by the object are no longer observed by the camera. This recursive co-location procedure takes about 10 minutes to perform.

After positioning the camera and SLM along the same optical line of sight, a perspective transform is calculated for the homography. Radial and tangential distortion by the camera lens is characterized by capturing an image of a checkerboard image and estimating the camera's intrinsic parameters and distortion coefficients. A homography is computed by first projecting a checkerboard pattern and capturing an image. The image is then undistorted and detected corner points are used to compute a perspective transform. After performing these calibration steps, the transformations are stored in look-up tables for later use and the system can be used anywhere without modification. These calibration steps were performed using functionality available in the OpenCV library.

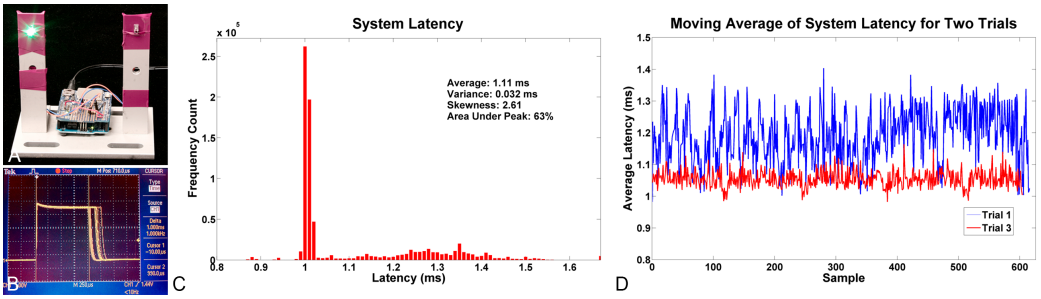

Figure 5.9: A: Circuit for measuring system latency consists of an LED and a phototransistor connected to a micro-controller board (Arduino Uno). To measure the system's reaction time, the micro-controller measures the time for the system to detect the illuminated LED then illuminates/dis-illuminates the phototransistor. B: Latency is observed on an oscilloscope (typical readout shown) and measured/recorded by the micro-controller board. C: Histogram (0.01 ms per bin) shows data collected over 30 minutes. The system has some uncertainty, but typically reacts within about 1 ms. D: Moving average (1 second intervals) of latency for a trial with and without uncertainty.

5.6 Measuring System Latency

As discussed in Section 5.4.2, the system is pipelined in three stages: image capture and transfer, image processing and transfer, and illumination. Latency of the system is the time between capturing an image and illuminating the scene. There are several factors that contribute to latency. Image size is directly related to camera/computer and computer/SLM transfer time, and image processing time. The size and number of detected objects also has an effect on latency requiring more processing time and thus increases latency. Lastly, the computer's operating system has timing jitter and interrupts that add uncertainty to the latency.

To measure system latency, a circuit was built to measure the system's time to react to an illuminated LED (Figure 5.9A). The circuit consists of an LED, phototransistor, and micro-controller (Arduino Uno). The high-level idea is to measure the response time of the system by enabling an LED and timing how long it takes for the system to detect the LED and project light onto the phototransistor. To achieve this, the system was programmed to illuminate the phototransistor every other frame. Observing the signal from the phototransistor with an oscilloscope reveals a step response as shown in Figure 5.9B. The plateau of the signal corresponds to the time that the phototransistor is illuminated. Time was measured with microsecond precision and recorded with the micro-controller.

Data were collected for thirty minutes evenly divided over six separate trials to assess repeatability. During these trials, the image resolution was 800×220, exposure time was 750 μs, and frame rate was 1 kHz. Latency for all the trials is shown in Figure 5.9C. Across the six trials, the system most often reacts within 1 ms and 63% of the time reacts within two standard deviations from the peak. The average reaction time for all six trials was 1.11 ms with a variance of 0.032 ms. The histogram also reveals uncertainty in the system. This variability was studied by averaging every 1 second worth of data. Shown in Figure 5.9D are averaged data for two trials: one trial with little variability and one trial with a lot of variability. The plot shows that, in either situation, latency consistently varies by small fluctuations within a narrow band. In the worst case, the fluctuations range from 1 to 1.4 ms.

Several strategies can be utilized to account for latency variability. The uncertainty can be simply included in the illumination pattern by artificially increasing the size of detected objects. Light throughput will decrease, but accuracy will improve. Alternatively, temporal information can be used to predict the location of detected objects. Care must be taken to ensure the prediction model does not add too much time to the system's latency. At high frame rates, a linear model

will suffice for most applications. The time to perform processing tasks was measured, in software, with a high resolution timer (Windows API). The average time for processing was 0.3 ms with a standard deviation of 0.04 ms (Figure 5.8). The time to send data to the DMD board over USB was measured for 5 minutes with an average of 0.76 ms and a standard deviation of 0.07 ms.

5.7 Programmable Headlight Prototype: Faster and More Compact

A second prototype was built using a camera for sensing, a computer for processing, and a custom-built high-speed projector for illumination. A diagram of our hardware implementation is shown in Fig. 5.10. This system is similar to that of [23] except with significant improvements to reduce system latency by at least 40% with minimized jitter.

5.7.1 Imaging

A computer with an Intel Core i7 3.6 GHz CPU interfaces between the camera and projector and processes the images. A monochrome, near-infrared camera (Basler) with a global shutter is used to capture images. The camera is capable of low latency and high data throughput because of its pipelined pixel architecture and extended CameraLink interface. The camera and projector are co-located following the procedure described in [23]. Camera exposure was synchronized to the projector to quantize system jitter in a deterministic manner. Synchronization also permits measuring system latency as explained in Section 5.8.1. A PCI express 2.0 frame grabber (Bitflow) is used to transfer data, without buffering, into DDR3 memory.

Software Architecture: The system's functionality is performed with multi-threaded software. The *acquisition thread* retrieves images from memory and copies them into an image buffer shared with a *processing thread*, where the image is processed (e.g., binarization, prediction, transformation, etc). After an illumination pattern is computed, it is stored in an image buffer shared with a *display thread*, which uses OpenGL to transfer the illumination pattern to the projector via the custom board described in Section 5.7.2. Linux Kernel 3.8 with RT Patch (3.8.14.15.rt+) was used to ensure deterministic scheduling performance and hyper-threading was disabled to eliminate cache-coherence issues resulting in

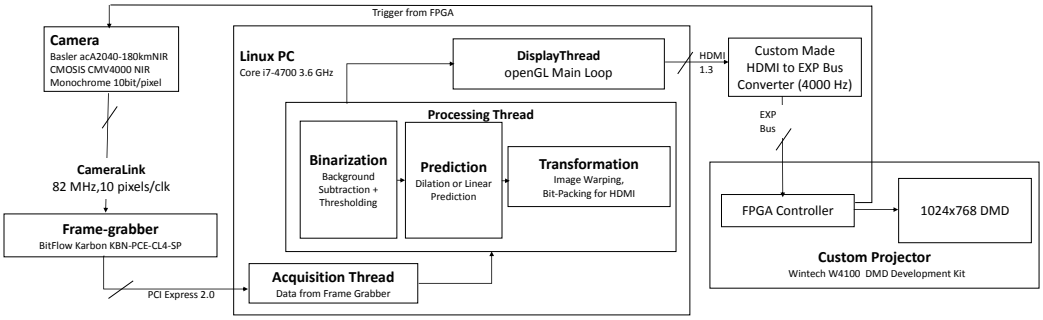

Figure 5.10: Diagram of our hardware implementation of a binary reactive visual system. Systems for other applications such as high dynamic range imaging and reconstruction will have similar implementations, but the same main components of sensing, processing, and projection.

minimized jitter.

Image Processing: The processing component consists of image binarization, object location prediction, image warping, and bit-packing. Quantitative evaluation of system performance is discussed in Section 5.8.2.

Binarization: Objects are segmented as follows:

$$B(x,y) = \begin{cases} 1 & \text{if } |I(x,y) - I_{bg}(x,y)| > \alpha \\ 0 & \text{otherwise,} \end{cases}$$

where α is an intensity threshold and $I_{bg}(x,y)$ is a pre-computed background image. The 8-bit image is subtracted from the background image and then the resulting image is thresholded. Since both of these are per-pixel operations, they were combined and written in AVX2 optimized SIMD vectorized code to increase computational efficiency.

Prediction: As discussed in Section 5.2.1, two approaches for estimating future locations of objects were investigated.

Image Warping: Image distortion caused by the camera lens and the transformation between the camera and projector planes are stored in a look-up table. The look-up table maps pixels from the projector's coordinate system to the camera's coordinate system to guarantee that every pixel has a correspondence. Access of the look-up table was written in AVX2 to increase the speed of memory operations. The result of the transformation is a binary image containing the illumination pattern sized to the projector resolution.

5.7.2 Projection

A high-speed projector (4,000 Hz) was built to achieve high-speed illumination. The projector consists of a DMD development board (WinTech W4100) [28] and a 4,800 Lumen light source. One of the bottlenecks identified in [23] was the time to transfer the illumination pattern to the projector over USB 2.0 ($0.6 - 0.7$ ms). To avoid this bottleneck, we designed a custom module to interface the PC using HDMI directly to the EXP bus interface on the DMD board. The HDMI board uses an integrated circuit (TI TFP401a) to parallelize the HDMI signals to a 28-bit video signal with a pixel clock rate of 177 MHz.

The binary illumination pattern is transferred over HDMI as a 24-bit image. Since the illumination pattern is binary, it is packed into a low-resolution, 24-bit image. The bit-packed image is transferred to the projector at a regular interval of 250 μs. The FPGA of the DMD board was programmed to expand the bit-packed image to the full resolution of the projector (1024×768). The board can transfer a full resolution image in 250 μs allowing a display rate of 4,000 Hz, which is 3 times faster than [23].

System	**Our system**		**[23]**	**[12]**
Resolution	960×340	960×170	1000×340	244×120
Exposure	$100\mu s$	$100\mu s$	$100\mu s$	$5000\mu s$
Image Transfer	$925\mu s$	$520\mu s$	$925\mu s$	$4200\mu s$
Acquisition Binarization Warping Bit-Packing†	$20\mu s$ $20\mu s$ $160\mu s$ $20\mu s$	$10\mu s$ $12\mu s$ $80\mu s$ $10\mu s$	$\geq 300\mu s$‡	$4100\mu s$‡
Display	$250\mu s$	$250\mu s$	$760\mu s$	$4200\mu s$
Total	$1495\mu s$	$992\mu s$	$\geq 2085\mu s$	$17500\mu s$

† Bit-packing not performed in [23] and [11]

‡ Processing time increases as the number of detected objects increases

Table 5.1: Latency of our system with two different image resolutions compared to other reported systems. The latency of our system is independent of the number of objects detected and can be used to emulate other systems by means of a software delay.

5.8 System Response Time

5.8.1 Measuring End-to-End Latency

To measure the reaction time of the system, the FPGA of the DMD board was programmed to output a trigger pulse when a new illumination pattern was received. Latency was measured with a logic analyzer (Saleae Logic 8) as the time between successive trigger pulses. Because the HDMI clock transfers data asynchronously at 4,000 Hz, jitter is quantized at 250 μs intervals. For example, if there is a delay in the system, (e.g., because of image analysis), the image will have to be displayed after a delay of 250 μs.

In addition to measuring round-trip latency of our system, we also measured the time to execute various processes. The latency of the system depends on various factors including image size (transfer and processing), algorithm complexity, and system uncertainty. Latency of our system for two resolutions (only binarization, warping, and bit-packing performed) are detailed in Table 5.1. The latency of similar systems [12, 23] are shown in the same table.

5.8.2 Effect of Prediction Algorithms on Latency

To compare system performance with different prediction algorithms, images generated from the simulator (Section 5.2) were used to guarantee data variability and repeatability for different trials. Latency was measured using the same method in Section 5.8.1. Data were collected for a minute for each trial and summary statistics were computed.

The image resolution for which the camera field of view covers the projector's field of view is 960×680. To decrease the system's reaction time, camera images were vertically decimated by a factor of both 2 and 4 yielding image resolutions of 960×340 and 960×170, respectively. Kernel sizes are reported as their effective size at full image resolution. For example, a kernel of radius 9 in the full resolution image has a radius of 5 for decimation by a factor of 2 and a radius of 3 for decimation by a factor of 4.

For objects moving in an arbitrary direction (like the exploding ping pong balls), dilation with an isotropic kernel is sufficient. However, for particles moving in a generally known direction (falling snow), it is a better strategy to dilate in the known direction. As shown in Fig. 5.11, increasing the kernel size significantly increases latency. Latency is reported as an average, minimum, and maximum value. Latency measured for different snowfall rates ranging from a light flurry

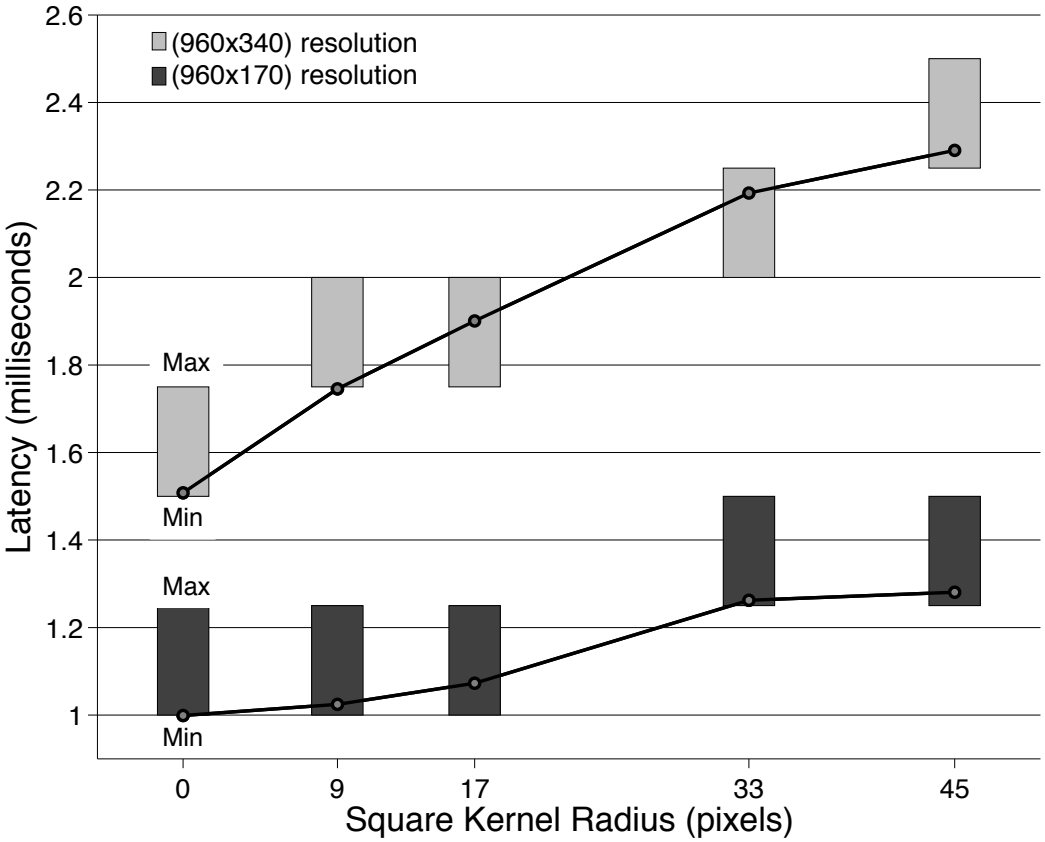

Figure 5.11: Increasing the size of the dilation kernel increases the latency of the system. Kernel radius is reported as the effective size at a full image resolution of 960×680. Circles indicate average latency, which are quantized by the 250 μs jitter of our system.

to a blizzard are shown in Fig. 5.12. Latency increases because the algorithm depends on the number of objects detected, which increases with the snowfall rate.

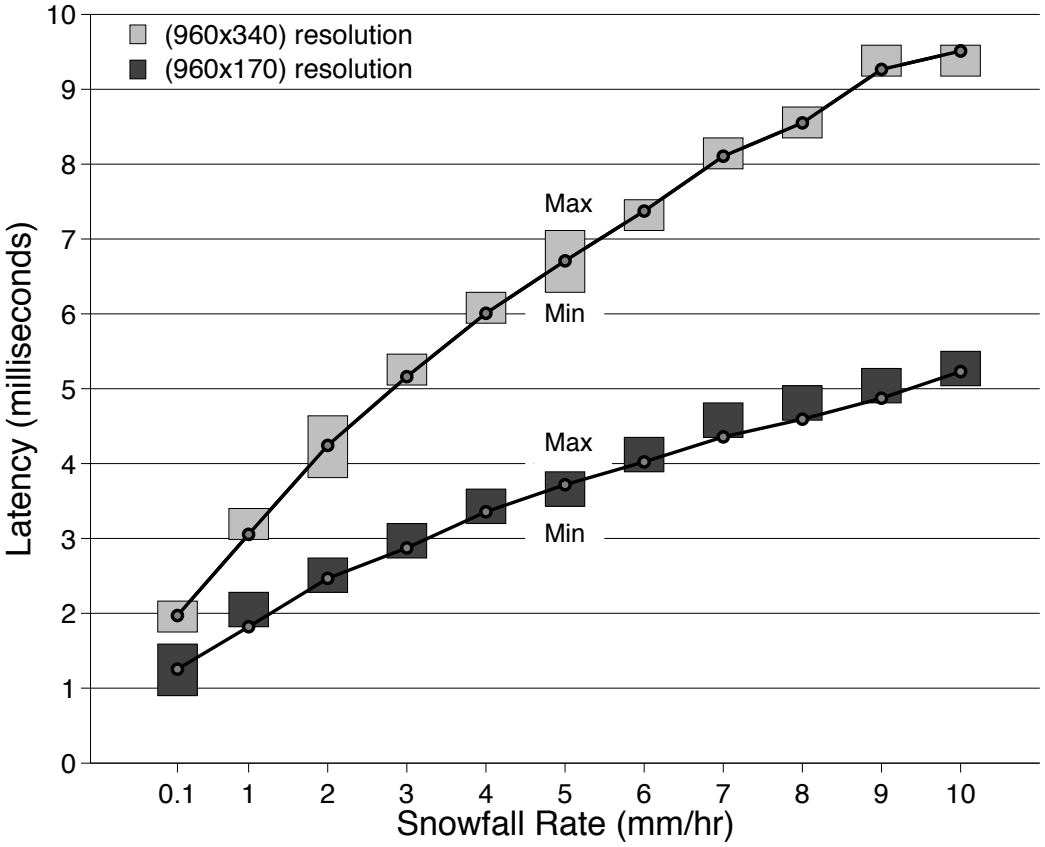

Figure 5.12: The latency of a system with linear prediction steadily increases with snowfall rate because our algorithm requires more computations to process additional detected snowflakes. Recall, that our system has 250 μs jitter, which is demonstrated by the average latency as indicated by circles.

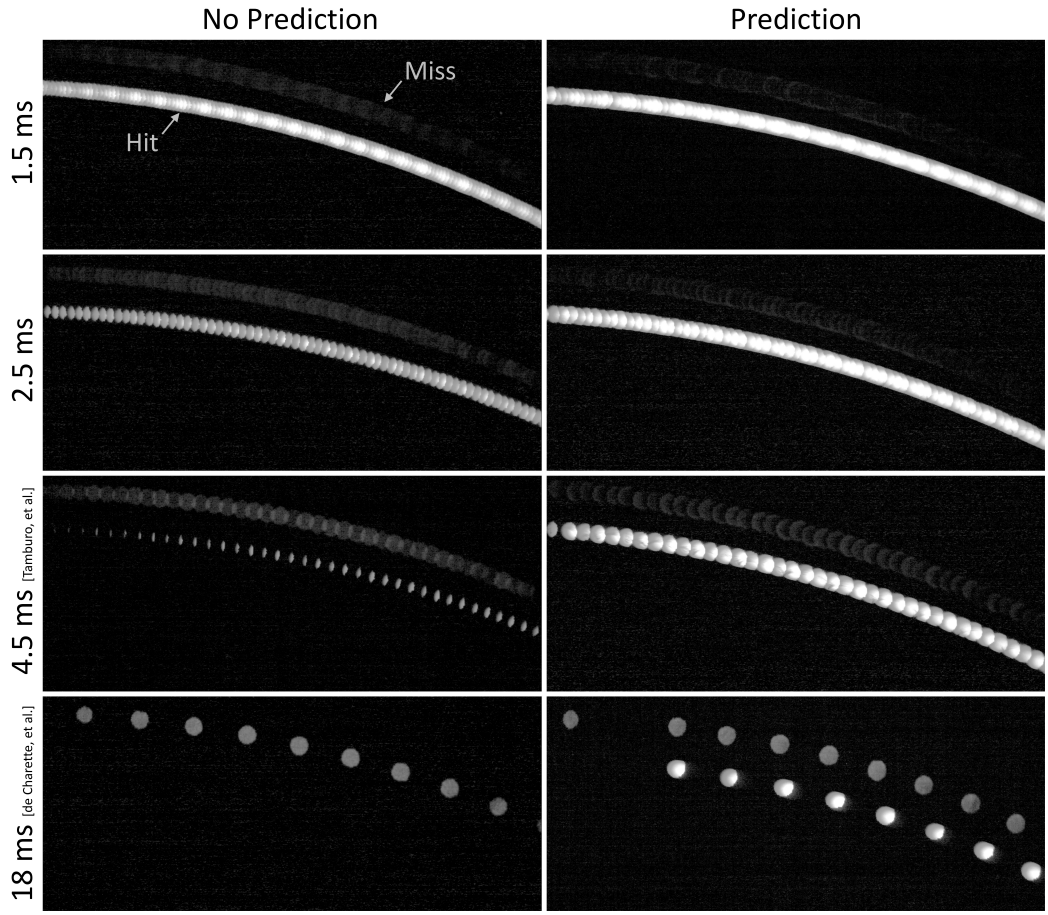

Figure 5.13: A ping pong ball with projectile motion was illuminated by the system while capturing a long exposure image. Bright streaks (hit) show correct illumination of the ball while dark silhouettes (miss) in the background show the error caused by a slow reaction of the system. Results from our system are in the top row. Results using latencies from [23] and [12] with our system are shown in the third and fourth rows, respectively.

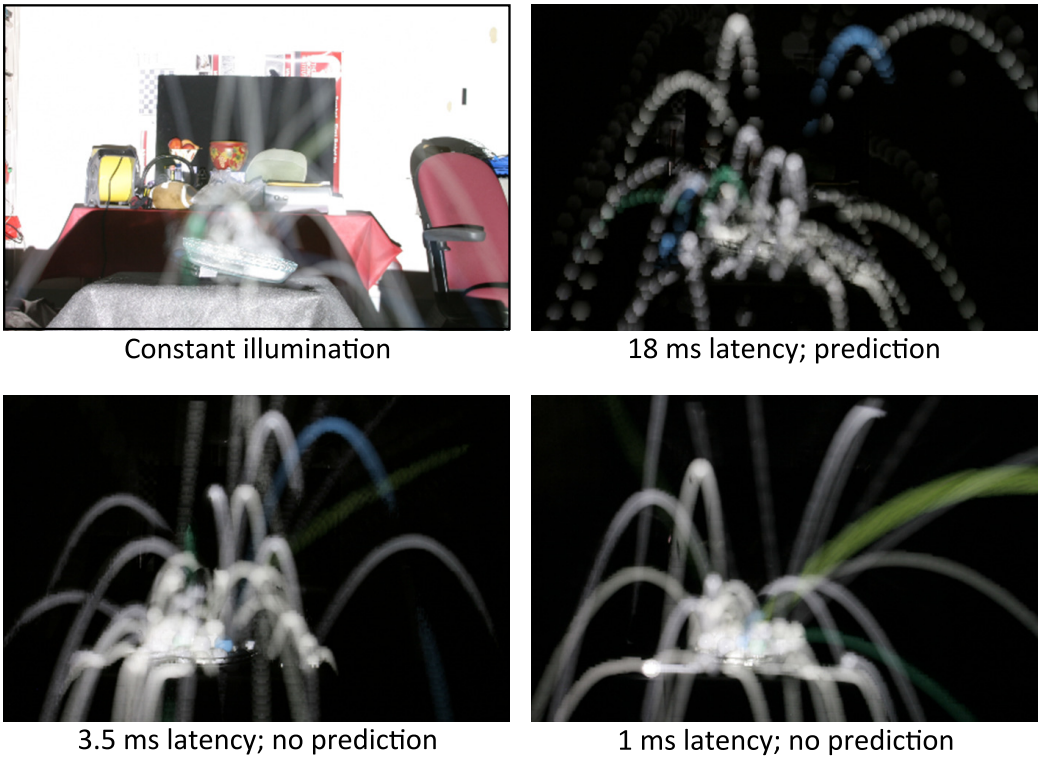

Figure 5.14: Visual evaluation of system performance for rigid objects with projectile motion. A tennis ball is thrown into a bowl of ping pong balls. The first image shows constant illumination. In subsequent images, ping pong balls are adaptively illuminated.

Figure 5.15: Visual evaluation of system performance for small objects moving chaotically. Artificial snowflakes are dropped. Still frames from a video captured at 30 fps are shown. In (a), snowflakes are constantly illuminated. In (b)-(d), snowflakes are adaptively dis-illuminated while increasing the size of the dilation kernel resulting in higher system accuracy and lower snowflake visibility.

Chapter 6

Findings

6.1 Glare Free High Beams

Glare from the headlights, especially high beams, of oncoming vehicles cause significant stress and distraction at best and temporary blindness at worst. Trucks and other vehicles with headlights at high positions are the worst offenders. Although, glare is not often reported as a cause of accidents, hundreds of fatal night crashes attribute glare as a contributing factor every year [6]. Glare is especially problematic for the elderly whom take eight times longer to recover from glare as compared to a 16-year old [13]. Although high beams are a nuisance to other drivers, they are beneficial on narrow, curvy, and poorly lit roads, especially in rural areas where wildlife routinely jumps onto the road.

Glare free headlights are currently being deployed by car companies, e.g., [27], [18], [16], [14]. The details of their systems are publicly unavailable, but it is known that these systems utilize multiple LEDs and sensors placed at different locations in the vehicle, e.g, [8], [26], [22]. Based on this information, it can be inferred that spatial resolution is limited to the number of LEDs. Camera frame rates of these headlight systems are limited to 30 - 60 Hz and thus have high latency [4], [1], [2]. In this section, it will be shown that a high-resolution SLM with low latency produces the best light throughput.

6.1.1 System Requirements and Comparisons

Computer simulations were performed to determine the latency required to maintain high light throughput. Camera parameters and the position of our prototype

on a vehicle were used in simulations where two vehicles traveled towards each other at 225 kph on a two-lane, straight road. Detection and prediction were set to be error-free guaranteeing that only system latency contributed to light throughput. Light throughput was calculated for latencies of 2, 16, 30, 50, and 100 ms (Figure 6.1A). Throughput remains above 90% for all latencies tested when the vehicles are farther than 20 m apart. The reason for this is the oncoming vehicle is moving towards the camera and its position in the image has little variation. However, as the vehicles move closer towards each other, light throughput substantially decreases with higher latency. It is clear that the system needs a latency of at least 2 ms to maintain 90% light throughput when the vehicles are close. The same would be true for vehicles in further lanes or on curved roads.

Computer simulations were conducted to compare the performance LED-based headlights to DMD-based headlights with the same latency. Since specific details of LED-based systems are publicly unavailable, several assumptions were made: (a) LEDs were positioned in a linear array parallel to the road and (b) all LEDs in the array that would illuminate the oncoming driver are disabled. Simulation results are shown in Figure 6.1B along with those of the DMD-based system. The low spatial resolution of LED-based systems results in lower light throughput and also creates flicker (abrupt changes in light throughput) for the driver. The flicker can be reduced by turning off more LEDs, but with the trade-off of sacrificing light throughput.

6.1.2 Our Headlight Design as Glare Free High Beams

The glare problem and our solution is illustrated in Figure 6.2A. Headlights from oncoming vehicles are detected in the captured image. Headlights are detected using the assumption that they are the brightest objects in the system's field of view. A very short exposure (100 μs) time is used and the image is thresholded. False detections can be reduced by excluding connected components that are too small to be headlights. Once the locations of the vehicles are known in the camera's reference frame, they are transformed to the headlight reference frame and the spatial light modulator blocks light in that direction. Since the resolution offered by the SLM is very high, only a small region above the detected headlight overlapping the oncoming driver's head is dis-illuminated. This type of beam blocking can be done for any number of oncoming drivers without significant loss of illumination. Compared to the system settings used to evaluate latency in Section 5.6, the image resolution was increased to 1000×340 to provide the largest field of view possible resulting in a system latency to 2.5 ms.

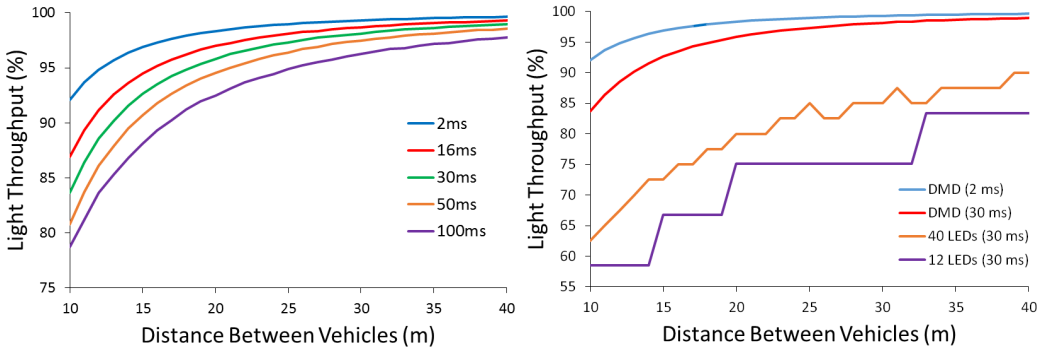

Figure 6.1: Results of computer simulations of glare free headlights. Detection and prediction are assumed to be perfect and vehicles were traveling towards each other in adjacent lanes at a relative speed of 225 kph. Left: Light throughput as a function of distance between vehicles for different system latencies. Right: Light throughput for DMD- and LED- based glare free headlights for different latencies. Simulations show lower latency and higher resolution results in higher light throughput and accuracy, which will be even more relevant for curvy, multilane roads with multiple vehicles.

6.1.3 Road Demonstrations

The system was tested on the road at night with three oncoming vehicles. Figures 6.2B-D show video frames captured from inside vehicles driving towards the programmable headlight. In Figure 6.2B, the blinding glare as the vehicles near each other is shown. Figure 6.2C and D show the benefit of our glare free headlight. Clearly, the difference in visibility is significant allowing drivers to see the road, vehicle, and surroundings. The prototype was able to function for all three drivers at the same time with little light loss. The average light throughput was calculated from saved images to be 93.8% with a standard deviation of 3.3%. In Figure 6.2F, tail lights were detected to avoid illuminating the driver's rear-view mirror and glaring them from behind. As shown in Figure 6.3, there is no discernible difference to the driver with the programmable headlight when the glare free function is enabled. The odd shape of the light beam is due to the system's position and the perspective of the capturing device. Installation in the headlight bay will create a more uniform shape and the spread of the light beam can be increased with a wide angle lens.

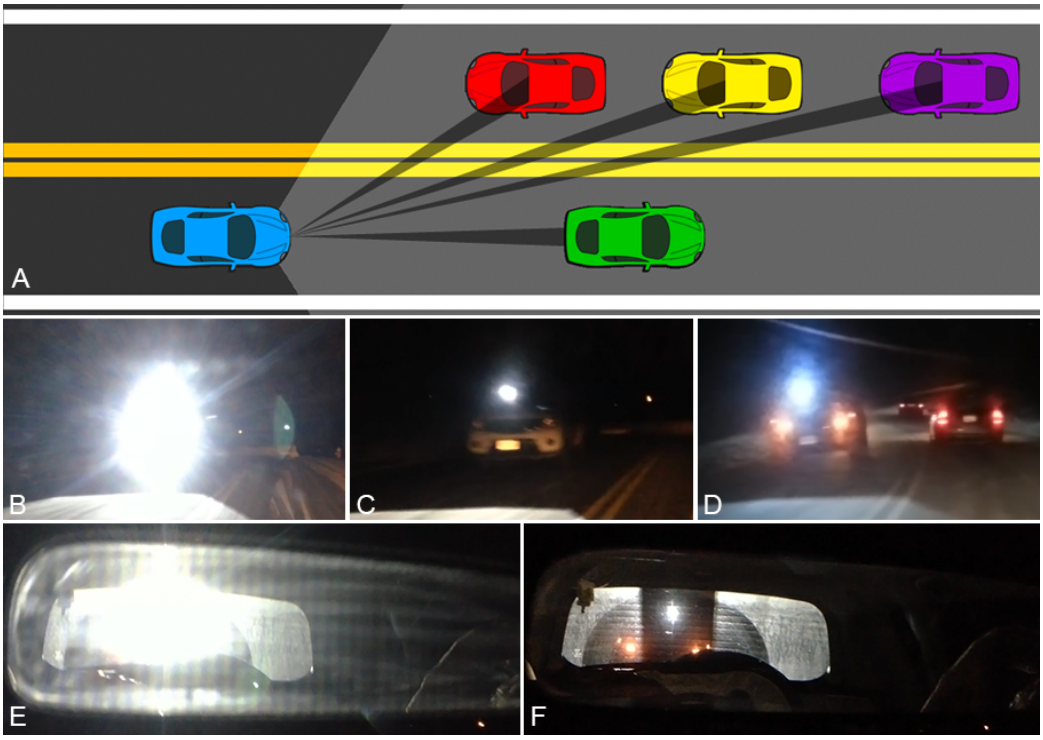

Figure 6.2: A: Illustration for eliminating high beam glare. Vehicles are identified and small regions around drivers are dis-illuminated while maintaining illumination elsewhere. Drivers with programmable headlights can then potentially use high beams without worry. Middle row shows view while driving towards our prototype. B: Glare typically seen from high beams (glare free feature disabled). C: Reduced glare when the glare free feature of our headlight is enabled. D: Glare free headlights allow the driver to better see other vehicles on the road. E: Glare in a rear view mirror caused by a following vehicle. F: Tail lights are detected to avoid illuminating the rear-view mirror.

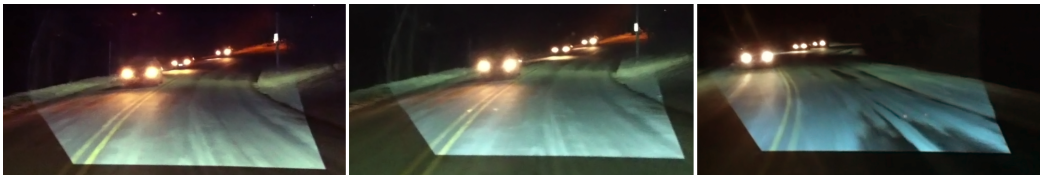

Figure 6.3: View shown from the perspective of the vehicle equipped with our prototype. Left: Glare free feature is disabled acting as a typical high beam. Middle and Right: Glare free feature is enabled detecting multiple oncoming vehicles and reducing light only in the direction of each driver. Notice no discernable difference between images.

6.2 Demonstrating Headlight Versatility

Thus far, computer simulations and demonstrations have shown that the proposed headlight design is advantageous to current glare free headlight designs. Our headlight can also be programmed to perform other tasks, whereas, other advanced lighting systems may require additional light sources, sensors, mechanical parts, etc., or are insufficient due to low spatial resolution or high latency. Here we show several tasks, such as visibility improvement in snowstorms (using artificial snow) and illuminating roads with better contrast and lane definition (visual warning of obstacles can be seen at [3]). Also shown is a computational photography application to examine high-speed events.

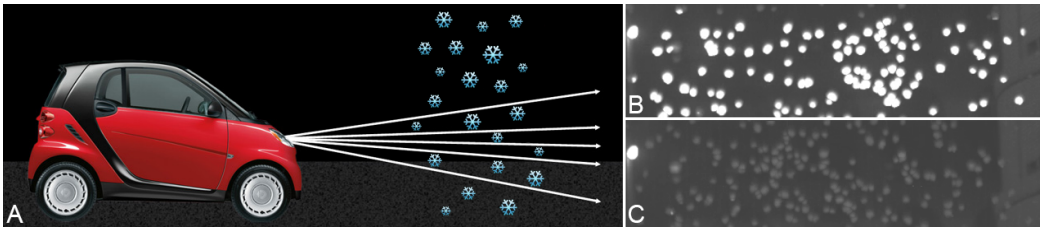

Figure 6.4: A: Our headlight has unprecedented resolution over space and time so that beams of light may be sent in between the falling snow. Illustration adapted from [12]. B: Artificial snowflakes brightly illuminated by standard headlight. C: Our system avoids illuminating snowflakes making them much less visible.

6.2.1 Improving Visibility During Snowstorms

Driving in a snowstorm at night is incredibly difficult and stressful. Snowflakes are illuminated brightly and distract the driver from observing the entire road. Researchers in computer vision have proposed methods for removing snow from videos [21], [29], [15]. Processed videos can be displayed for the driver, but current implementations are not intuitive and, at times, distracting for the driver. We can address this problem with a solution similar to that for glare free, i.e., reacting to detected bright objects. The main difference, however, is that the density, size, and speed of snowflakes requires high-resolution, low-latency illumination to be effective. Therefore, we exploit the high-resolution and fast illumination beam control of our prototype to distribute light between falling snowflakes to reduce backscatter directly in the driver's visual field (Figure 6.4). However, this application is significantly more challenging since (a) the size of snowflakes is very small compared to an easily detectable vehicle and (b) the quantity of snowflakes is several orders higher than the number of cars on the road. The goal is to send as much light as possible from the headlight to sufficiently illuminate the road for the driver while dis-illuminating snowflakes.

Computer simulations performed in [12] demonstrate that the idea is feasible. They estimate that, for a vehicle traveling at 30 kph, the system's latency needs to be 1.5 ms or less to have high light throughput and accuracy. We demonstrate improved visibility outside at night with artificial snowflakes. Snowflakes were detected by performing background subtraction and binary thresholding. To compensate for any small detection errors, dilation with a structuring element of a radius equivalent to that of a snowflake was applied. The visibility improvement can be seen by comparing Figures 6.4B and 6.4C. Even though the snowflakes fall chaotically, no prediction was required because of the system's fast speed. For comparison, the system by [12] (13 ms latency) was demonstrated for rain drops falling along a straight path and required a linear prediction model.

The effect of increasing the size of the dilation kernel is shown in Fig. 6.5. Artificial snowflakes (styrofoam beads) were dropped while being dis-illuminated by the system. As the kernel size is increased, accuracy increases, and consequently, visibility of the snowflakes decreases. Larger kernel sizes decrease the visibility of snowflakes while increasing false positive error - evident by larger dark streaks visible on the road. The requirements of this trade-off space will vary between applications. Results of these visual evaluation experiments confirm those of the simulations.

Figure 6.5: Visual evaluation of system performance for small objects moving chaotically. Artificial snowflakes are dropped. Still frames from a video captured at 30 fps are shown. In (a), snowflakes are constantly illuminated. In (b)-(d), snowflakes are adaptively dis-illuminated while increasing the size of the dilation kernel resulting in higher system accuracy and lower snowflake visibility.

6.2.2 Improved Lane Illumination

Sometimes the road is not clearly visible and no amount of illumination from a standard headlight can assist the driver. A few examples of such situations are snow covered roads, roads without lane markings or shoulders, and poorly lit roads. Our prototype can be used to brightly illuminate only the driver's lane to provide them with a visual guide. Opposing lanes, curbs, and sidewalks can be dimly illuminated to create a strong contrast with the driver's lane and also provide sufficient illumination to see obstacles (Figure 6.6A). For this application, images do not need to be captured or analyzed, and objects do not need to be tracked.

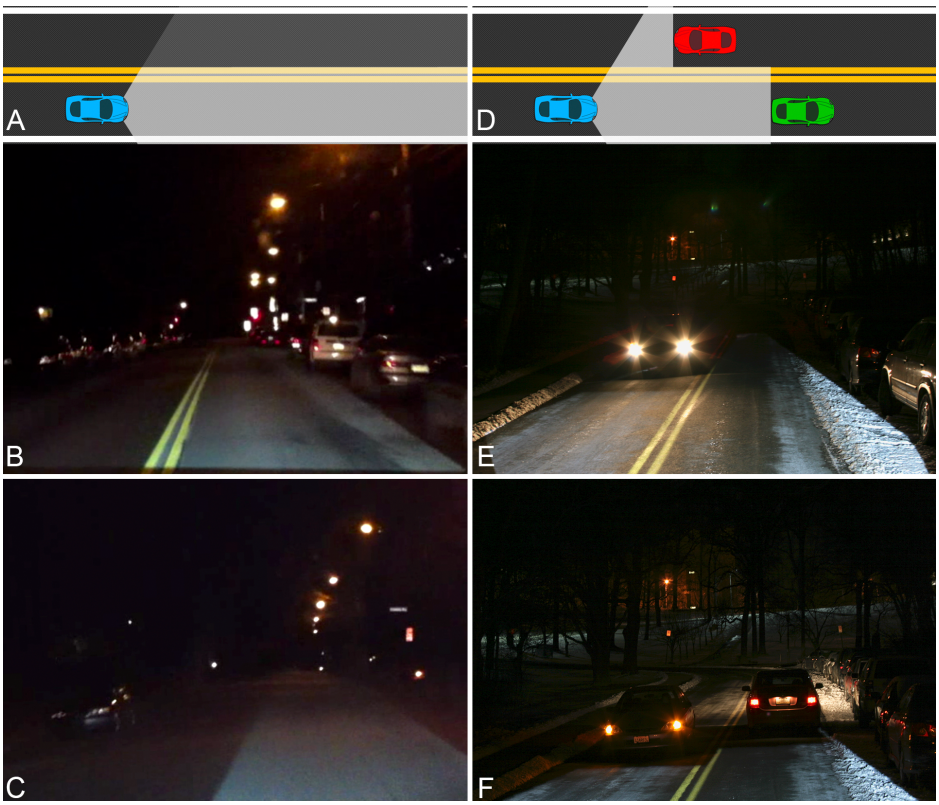

Figure 6.6: A: Concept of illuminating the driver's lane with high-intensity light and illuminating the adjacent lane dimly to improve the contrast of the driver's lane. B: Driver's lane more brightly illuminated than the adjacent lane. C: Demonstration while driving on an unmarked road. D: Concept of adjusting lane illumination based on the presence of other vehicles. E: Illumination for the left lane stops at the oncoming driver to avoid projecting lane patterns on the vehicle. F: Lane illumination stops in front of the vehicle in the adjacent lane and behind the vehicle in the driver's lane.

After computing the homography with the road plane, the headlight acts only as an illumination device. For proof-of-concept, illumination patterns were predetermined for the stretch of road where experience were conducted. In practice, the position and speed of the vehicle will be used to dynamically determine the illumination patterns required for the road.

In Figure 6.6B, the driver's lane and lane markings are fully illuminated, and the adjacent lane is dimly illuminated. The same contrast is used while driving

on a dark, unmarked road in Figure 6.6C. The opposing lane is dimly illuminated while the driver's lane remains fully illuminated creating a demarcation line for the driver to follow. Vehicles driving on the illuminated lane will experience disorienting illumination patterns because the system is calibrated to illuminate the road plane. Therefore, the beam can be adjusted where vehicles are detected in either lane as illustrated in Figure 6.6D. The adjacent lane can be illuminated up to the location of an oncoming vehicle while maintaining full illumination of the driver's own lane (Figure 6.6E). Illumination can be controlled in the presence of vehicles in both lanes as well (Figure 6.6F).

6.2.3 Early Visual Warning of Obstacles

Similar to improving visibility of the road and driving lanes, the headlight can be used to improve the visibility of obstacles on the road. Driving long distances or the same route routinely reduces a person's attention to obstacles in their peripheral vision. Coupled with low peripheral illumination by headlights and by the environment, it is not uncommon for people to slowly visually identify obstacles on the side of the road - often times too slow to avoid collision. Obstacles of interest could include pedestrians, wildlife, bicyclists, construction workers, traffic signs/barriers, etc.

Our simulations with the fast moving rigid objects indicate that our prototype system is capable of tracking objects moving up to 130 kmph across the camera's field of view. In the first image of Fig. 6.7, the bicyclist is barely visible even with street lighting and the vehicle's high beams. As the bicyclist crosses the road, our headlight detects him and illuminates the bicycle as well as the road in the area of the bicycle. The bicyclist was intentionally not spotlighted to avoid glaring him.

Figure 6.7: Example of spotlighting a bicyclist to provide an early visual warning for driver's on the road. In the first image, the bicyclist is difficult to see even with the high beams on. In subsequent images, the bicyclist crosses the road as our headlight detects it and illuminates the road in the direction of the bicyclist and the bicycle itself. The bicyclist is intentionally not spotlighted to avoid glaring him.

Chapter 7

Conclusions

The headlight should not be a passive device that can just be completely switched on or off. It should be capable of adapting to the environment. Moreover, the design for adaptive headlights should not be a one-off solution. It should be programmable capable of performing many different tasks to help the driver in various road environments. Our headlight design provides unprecedented light beam control over angle and time. Essentially, the full headlight beam can be split into hundreds of thousands of tiny little beams that can be turned on or off for very short durations (milliseconds). The flexibility and control of the headlight will allow us to perform numerous tasks for the first time: Allowing drivers to use high-beams without glaring any other driver on the road, allowing drivers to see better in snow, and allowing better illumination of road lanes, sidewalks and dividers. The prototypes we have built can react to the road environment within 1.0 to 1.5 milliseconds with a refresh rate of up to 4 kHz. This refresh rate will not cause any flicker to be seen by the human eye.

Chapter 8

Recommendations

Feasibility studies and road demonstrations with our adaptive headlight prototypes show that they can improve and enhance visibility for the driver at night during a variety of road conditions. We have shown that glare free high beams do not affect oncoming drivers. Unlike existing 'glare-free' headlights that automatically switch to low beams our design provides high light throughput (high beams) for the driver. Europe, unlike the U.S., permits a portion of the high-beam to be deactivated such as systems with matrix-LED headlights. Our demonstrations also show potential safety benefits by avoiding backscatter during rain/snow storms, spotlighting pedestrians, bicyclists, and animals, and enabling contrasting illumination of the driving lane. As car companies and headlight manufacturers develop adaptive headlight solutions, we recommend the U.S. Department of Transportation:

- Enact suitable regulations are in place to permit the deployment of adaptive headlights on rural and urban roads.
- Establish a set of standards for which adaptive headlights should meet in terms of performance, reliability, and durability.
- Approve and authorize a rating system for which adaptive headlights can be graded based on their safety impact.

Bibliography

[1] Advanced driving assistance and active safety systems. *http://media.opel.com/media/intl/en/opel/vehicles/opel_eye/2009.html*.

[2] Hella group website. *http://www.hella.com/hella-com*.

[3] Illumination and imaging laboratory project web page for smart headlights. *http://cs.cmu.edu/∼ILIM/SmartHeadlight*.

[4] Mobileeye camera matrix. *http://www.mobileye.com/technology/development-evaluation-platforms/cameras*.

[5] N. H. T. S. Administration. Report on drivers' perceptions of headlight glare from oncoming and following vehicles. 2003.

[6] N. H. T. S. Administration. Nighttime glare and driving performance. 2007.

[7] N. H. T. S. Administration. Traffic safety facts 2013: A compilation of motor vehicle crash data from the fatality analysis reporting system and the general estimates system. 2013.

[8] M. Benz. Mercedes-benz announces new active multibeam led headlights. *Press Release*, 2013.

[9] C. Bölling. Osram presents future technologies for car headlights. *Osram Press Release*, 2013.

[10] F. M. Company. Next generation of ford motor company's headlights make nighttime driving safer. *Ford Press Release*, 2005.

[11] R. de Charette, R. Tamburo, P. Barnum, A. Rowe, T. Kanade, and S. Narasimhan. Fast reactive control for illumination through rain and snow. pages 1–10, 2012.

[12] R. de Charette, R. Tamburo, P. C. Barnum, A. Rowe, T. Kanade, and S. G. Narasimhan. Fast reactive control for illumination through rain and snow. *In: IEEE International Conference on Computational Photography (ICCP), Seattle, Washington*, 2012.

[13] A. F. for Traffic Safety. How to avoid headlight glare. 2013.

[14] P. Fröberg. Volvo cars makes driving at night safer and more comfortable with innovative, permanent high beam. *Volvo Car Group Press Release*, 2013.

[15] K. Garg and S. K. Nayar. Detection and removal of rain from videos. *In: IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR), vol. I, pp. 528-535*, 2004.

[16] N. Giesen. New generation cls with the future's high-resolution precision led technology: Leading the way with better light. *Daimler Group Press Release*, 2014.

[17] V. Group. To the point: The new polo gti: Extremely strong and exceptionally fuel efficient. *VW Press Release*, 2010.

[18] H. Inc. Hella develops unique matrix led headlamp system with audi. *Press Release*, 2014.

[19] L. Rice. Headlight with single led module. *SAE Technical Paper 2010-01-0295*, 2010.

[20] D. Schuellerman. Ge lighting unveils high-performance headlamp lighting solutions. *GE Lighting Press Release*, 2012.

[21] Y. Shen, L. Ma, H. Liu, Y. Bao, and Z. Chen. Detecting and extracting natural snow from videos. *In: Information Processing Letters, vol. 110, pp. 1124-1130*, 2010.

[22] T. Söllner. Audi - the leading brand in lighting technology. *Audi Press Release*, 2013.

[23] R. Tamburo, E. Nurvitadhi, A. Chugh, M. Chen, A. Rowe, T. Kanade, and S. G. Narasimhan. Programmable automotive headlights. In *Computer Vision - ECCV 2014*, volume 8692 of *Lecture Notes in Computer Science*, 2014.

[24] A. Toshiyuki, K. Osamura, and M. Fujisawa. Controlled illumination for the object recognition with projector camera feedback. *In: IAPR Conference on Machine Vision Applications, pp. 152-155*, 2011.

[25] O. Wang, M. Fuchs, C. Fuchs, J. Davis, H.-P. Seidel, and H. P. A. Lensch. A context-aware light source. *In: IEEE International Conference on Computational Photography (ICCP), Cambridge, MA*, 2010.

[26] M. Wiese. Bmw innovations in vehicle lights. "dynamic light spot" for actively illuminating persons, the "glare-free high beam assistant" and full-led headlights provide even more safety at night. *BMW Press Release*, 2011.

[27] M. Wiese. Bmw lights the way into the future. *BMW Press Release*, 2014.

[28] WinTech. DLP discovery W4100 kit. 2014. `http://wintechdigital.com/Upfiles/W4100_Brochure.pdf`.

[29] C. Zhen and S. Jihong. A new algorithm of rain (snow) removal in video. *In: Journal of Multimedia, vol. 8, no. 2*, 2013.