ITS IDEA PROGRAM COMMITTEE

Roy Allen
Association of American Railroads

Clifford Bragdon
Dowling College

Daniel Brand
Charles River Associates, Inc.

A. Ray Chamberlain
Parsons Brinckerhoff

John German
City of San Antonio

Edwin Hauser
Kimley-Horn Associates

Michael Kushner
NextBus Information Systems, Inc.

Marshall Lih
National Science Foundation

George Parker
Assn. of International Automobile Manufacturers

Ray Pethtel
Virginia Tech

Noah Rifkin
Veridian Engineering, Transportation Group

Charles Wallace
University of Florida

Carol A. Zimmerman
Battelle

Liaison Members

Fenton Carey
RSPA

John Collins
ITS America

Christine M. Johnson
US DOT

Dennis C. Judycki
FHWA

Walter Kulyk
FTA

Robert McCown
FRA

Raymond P. Owings
NHTSA

Jeffrey F. Paniati
ITS Joint Program Office

Raymond E. Starsman
ITS America

Michael F. Trentacoste
FHWA

DOT Staff

Pamela Crenshaw

FHWA Staff

Robert Ferlis
David Gibson

FRA Staff

Lang Nguyen
Steve Sill

ITS America Staff

Steve Keppler

NHTSA Staff

Duane Perrin

IDEA Program Staff

Robert E. Skinner, Jr., Executive Director
Neil Hawks, Director, Special Programs Division
Keith Gates, ITS-IDEA Program Officer
Inam Jawed, NCHRP-IDEA Program Officer
Harvey Berlin, Transit-IDEA Program Officer
Charles Taylor, HSR-IDEA Program Officer
Richard Cunard, Division A Liaison
Debbie Irvin, Project Assistant

Mailing Address:
2101 Constitution Avenue, N.W.
Washington, D.C. 20418

Office Location:
GR 213
2001 Wisconsin Avenue, N.W.
Washington, D.C. 20007
202-334-3568

Innovations Deserving Exploratory Analysis Program

Final Report for IDEA Project ITS-61

IRIS: INTELLIGENT RANGING WITH INFRARED SENSORS

Principal Investigators: I. Kanellakopoulos and O. M. Stafsudd
Researchers: P. Nelson, A. Nelsen, R. Grundy-Valenzuela, Y. Tan, N. Cohen, and J. Zhao
UCLA Electrical Engineering Department

Transportation Research Board
National Research Council

2000

CONTENTS

EXECUTIVE SUMMARY

The objective of this ITS-IDEA project was the development of a new technology for ranging sensors called IRIS. This technology has the potential to produce a new generation of sensors with significantly lower cost and substantially increased accuracy and reliability than any of the technologies available today. Such sensors would be useful in almost every aspect of Intelligent Transportation Systems that requires knowledge of the position of the host vehicle relative to the vehicles around it and with respect to the roadway itself. This specific project focused on the simplest version of this sensor, called IRIS-1, which computes the distance relative to the preceding vehicle. The main accomplishment was the construction of a working prototype of the IRIS-1 sensor that was used to successfully perform a wide array of static and dynamic in-vehicle experiments.

The hardware part of the prototype includes

- a low-power eye-safe infrared **laser illuminator**,
- a custom-built **CCD receiver camera** based on the Texas Instruments TC 245 chip,
- a **mask emulator** for the expose-shift-expose process built by our team at UCLA using a Ronchi grating and an optical setup for the emulation,
- a fairly narrow **bandpass filter** centered at the frequency of the infrared laser illuminator, and
- a **laptop computer** that interfaces through its parallel port with both the illuminator and the receiver.

The software for our prototype was developed entirely by our team at UCLA; it runs on a laptop computer and consists of four basic modules:

- a **driver module** that activates the illuminator and the receiver in the synchronous fashion necessary for the expose-shift-expose process,
- an **image acquisition module** that transfers the image data from the CCD chip to the laptop,
- an **object detection module** that detects all reflective surfaces in the image and computes their position and size, and
- a **ranging module** that computes the relative distance to the preceding vehicle, assuming that the two brightest reflections in the image are the two taillights of that vehicle.

We used this prototype to conduct static experiments in our laboratory, as well as dynamic experiments with two vehicles in motion. In the static experiments, we ran our camera in high-resolution mode. This resulted in an update frequency of 1/3 Hz (one update every 3 s), and errors of ±1% for distances from 5 ft to 60 ft. For the dynamic experiments, we ran the camera in low-resolution mode, in which we read clusters of 6 pixels as a single one. This allowed us to increase the refresh rate by a factor of 6 to 2 Hz (one update every 0.5 s), but the errors also increased to ±2% for distances from about 6 ft to about 50 ft. To the best of our knowledge, these errors are lower than those achieved by any other ranging technology currently available. Furthermore, our experiments clearly indicated that these errors were mostly due

to the imperfections of our experimental setup, and would be significantly reduced in a commercially produced version of our sensor system.

These results prove that the IRIS technology works, and that it provides a clear path to a new generation of cheaper and better ranging sensors. A great deal of additional research and development are needed to realize the full potential of this idea, and now we know that this effort is fully justified and should be vigorously pursued.

1. GENERAL DESCRIPTION

1.1. Introduction

Intelligent Transportation Systems (ITS) technology is progressing at an ever-increasing rate, with exciting developments on all fronts—from driver information and assistance systems that enhance safety and comfort, to Automated Highway Systems (AHS) aimed at increasing traffic throughput, further enhancing safety, and reducing emissions. One of the main points of contention in the ITS community is the level of vehicle-to-vehicle and vehicle-to-roadway cooperation that can be assumed. In that respect, systems currently in various stages of research and development can be classified into three categories:

- **Autonomous systems** depend only on information obtained by the sensors located on the vehicle itself, usually relative distance and velocity to stationary objects and moving vehicles. They are therefore implementable in the immediate future, and in fact have started to appear as commercial products (collision warning, adaptive cruise control).

- **Cooperative systems** add information transmitted by neighboring vehicles, usually acceleration and steering inputs. Hence, they can perform more demanding tasks than autonomous systems, such as coordinated driving in a group, but their time to commercialization is likely to be longer.

- **Automated highway systems** add information obtained from the roadway infrastructure, such as messages regarding traffic conditions and road geometry, and lateral information from magnetic nails or reflective guardrails installed on the highway. Such systems can perform even more demanding tasks, like fully automated driving in a platoon, but must face many more obstacles (standardization, liability issues, public acceptance) on their way to implementation.

Commercial vehicles, in particular, will reap significant benefits from all stages of automation. Collision warning systems increase safety and reduce accidents, while adaptive cruise control enhances driver comfort and reduces fuel consumption and emissions. Fleet operators can further reduce their costs using cooperative scenarios like the *electronic towbar*, in which one manually driven vehicle is followed by two or three driverless automated vehicles. Finally, in automated highway systems, fully automated vehicles will be able to carry freight and passengers with significantly enhanced safety, increased fuel efficiency, and much more predictable travel times.

A critical automation component in any of these scenarios, especially for commercial vehicles, is the ranging sensor. Existing sensor technologies can only provide the reliability required for these applications by increasing the cost to levels beyond those acceptable to their potential customers. The TRB-IDEA project ITS-61, entitled *IRIS: Intelligent Ranging with Infrared Sensors*, which has now been successfully completed, addressed this issue by providing the necessary "proof of concept" for a new sensor technology, whose main advantages are its low cost, its high reliability, and its usefulness in all stages of ITS deployment.

1.2. Operating Principle

The IRIS system is composed of a low-power infrared diode laser illuminator, similar to those found in commercial compact disc players, with a relatively wide beam that is pulsed on/off, and a receiver consisting of a charged-coupled device (CCD) image sensor, similar to those used in inexpensive surveillance cameras, behind a narrow bandpass filter.

As shown in Figure 1, the system records two images of the preceding vehicle in rapid succession: the first one with the illuminating laser turned off, and the second one with the laser on. Due to the micrometer wavelength of the illuminator, almost all surfaces of the preceding vehicle and the surrounding objects scatter the incident laser beam and reflect only a negligible fraction

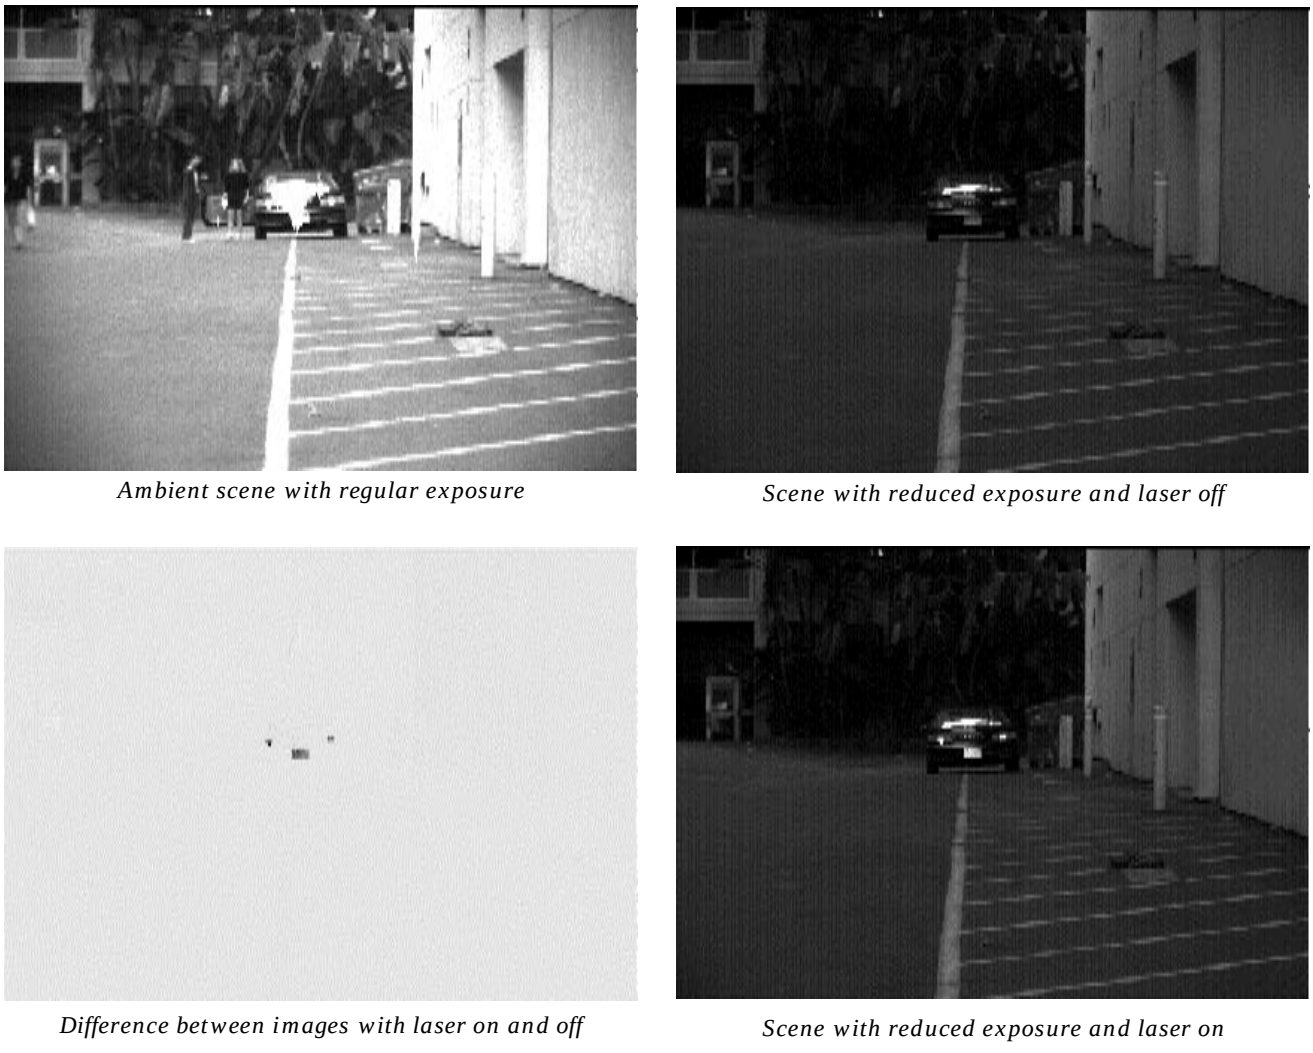

Ambient scene with regular exposure

Scene with reduced exposure and laser off

Difference between images with laser on and off

Scene with reduced exposure and laser on

Figure 1

IRIS operating principle: the original scene is reduced to a three-spot pattern, which is clearly distinguishable above the clutter.

back to the receiver to be recorded on the second (laser on) image. There are, however, some surfaces that reflect most of the incident laser beam back to its source: These are the passive reflective patches molded into the taillights of all modern vehicles, as well the reflective paint used in the license plates of many states, which act as efficient retroreflectors in the illuminator's wavelength. Subtraction of the first (laser off) image from the second (laser on) one yields a clear sharp image of these reflective surfaces against a flat background.

It is particularly important to note that, even in this simple example, the clutter is essentially nonexistent. This is due to the combined effect of

- the narrow bandpass filter, which significantly reduces the contribution of broadband sources such as the sun or artificial lights without affecting the laser returns, and

- the image subtraction process, which eliminates virtually all returns except those produced by laser illumination.

With the modulated illuminator and synchronous detection of the IRIS system, the retroreflectors and license plate are the dominant features in the subtracted image. The IRIS computational requirements are thus significantly lower than those of a conventional approach, which would use two-dimensional image processing techniques to extract the position of the vehicle from the cluttered scene of the original image. The distance to the preceding vehicle is then easily computed from the resulting clear pattern via standard triangulation schemes.

1.3. Advantages and Limitations

The IRIS sensor technology features several key properties, which can be viewed as significant advantages over existing technologies:

- As seen above, the signal-to-noise ratio is very high; this guarantees high reliability.

- The low cost of the components (a laser illuminator similar to those used in CD players and a CCD chip similar to those used in inexpensive surveillance cameras), combined with the simplicity of the concept, results in a much lower total system cost than any of the currently available technologies.

- Because every vehicle on the road has reflective taillights and many have reflective license plates, there is no need for installation of additional infrastructure.

- Due to the ability of the retroreflectors in taillights to reflect incident radiation from a wide range of angles back to its source, the IRIS sensor can see vehicles at angles up to 40 degrees on either side; its total field of view is, therefore, on the order of 80 degrees, which is wider than existing sensors.

- There is no need for on-line pointing or focusing; this eliminates any moving parts, hence increasing the robustness of the system while decreasing its cost.

- Ambient light conditions are irrelevant to the sensing process.

- Mechanical vibrations generated by the engine or the road do not affect the sensing process; hence, no active vibration cancellation is required.

Thanks to these features, the IRIS sensors have immediate applications in autonomous ranging for collision warning and adaptive cruise control systems currently under development by several automobile manufacturers, as well as the potential for significant long-term impact in AHS deployment, through their use for autonomous and cooperative ranging and even for vehicle-to-vehicle communication. In order to appreciate the full potential of the IRIS system, it is important to understand what its capabilities and limitations are.

Lateral and longitudinal ranging. When IRIS is locked on a vehicle, it monitors the apparent position of its taillights and license plate on the CCD chip. This provides enough information to measure the relative longitudinal distance, lateral deviation, and yaw angle between the leading and the following vehicle. However, in order to compute these quantities, the IRIS sensor needs to know a baseline lateral distance, such as the distance between the taillights of the preceding vehicle. This is due to the fact that IRIS uses triangulation schemes, so the length of at least one of the sides of the triangle needs to be known. This is not a problem in an experimental setup, where the distance between the taillights and license plate of the preceding vehicle will be exactly known. However, in a general highway environment, the IRIS configuration described thus far cannot yield accurate absolute measurements, since the distance between taillights varies from one vehicle to the next.

This problem can be eliminated through a straightforward modification of the sensor configuration: instead of using one CCD receiver, one can use two that are mounted at a known lateral distance from each other on the host vehicle. To distinguish between these two configurations, we will denote them as **IRIS-1** and **IRIS-2** depending on whether one or two receivers are used. The IRIS-2 sensor uses the distance between its receivers as the baseline, and then measures the apparent positions of the same object (such as left or right taillight) on each of its two receivers to implement its triangulation scheme. Thus, IRIS-2 can compute absolute distance at the additional cost of a second receiver and a more elaborate vehicle installation procedure. Although IRIS-1 cannot measure absolute distance, it can easily lock onto a vehicle in front and track **changes** in the distance with high reliability, since these changes are measured as changes in the apparent separation of the taillights on the single CCD receiver.

To place the above discussion in a more precise framework, let us first consider the IRIS-1 sensor: it calculates the lateral and longitudinal distances as well as the yaw of the preceding vehicle through a straightforward computational scheme. The only data necessary for this scheme to be carried out are: the focal length of the sensor (f), the taillight separation of the preceding vehicle (d) (assuming that the license plate is positioned exactly in the middle of that distance), and the location of the recorded reflections of these surfaces on the CCD image (u_l, u_c, u_r). The system setup is described in Figure 2.

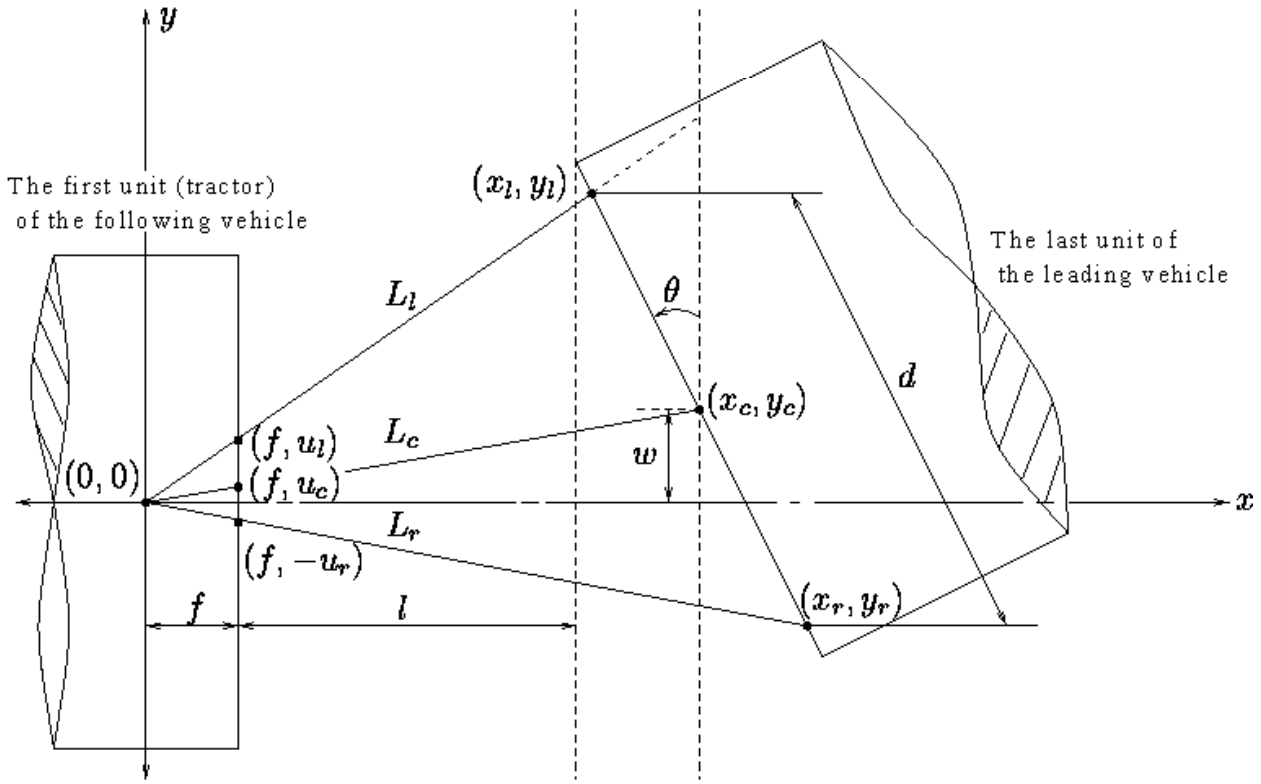

Figure 2

Geometry for extracting lateral and longitudinal information from IRIS-1.

From the CCD recorded information we define the following line functions:

$$L_l: y = \frac{u_l}{f} x$$

$$L_c: y = \frac{u_c}{f} x$$

$$L_r: y = -\frac{u_r}{f} x$$

These lines intersect with the leading vehicles' taillights and license plates at the points (x_l, y_l), (x_r, y_r), and (x_c, y_c) respectively. Applying this information as described in Figure 2 we arrive at the following equations:

$$\left.\begin{aligned} d &= \sqrt{\left(x_l - x_r\right)^2 + \left(y_l - y_r\right)^2} \\ x_c &= \frac{x_l + x_r}{2} \\ y_c &= \frac{y_l + y_r}{2} \\ y_l &= \frac{u_l}{f} x_l \\ y_c &= \frac{u_c}{f} x_c \\ y_r &= -\frac{u_r}{f} x_r \end{aligned}\right\} \begin{aligned} \theta &= \tan^{-1}\left(\frac{x_l - x_c}{y_l - y_c}\right) \\ l &= \min\left(x_l, x_r\right) \\ w &= y_c \end{aligned}$$

This set of six equations can be solved to acquire the relative yaw θ, the relative lateral deviation w, and the relative longitudinal distance l between the leading and the following vehicle. However, the IRIS-1 scheme clearly needs prior information about the taillight/license-plate geometry of the preceding vehicle, while the use of two CCD sensors installed on the front of the following vehicle eliminates that need. In fact, as we will see below, the IRIS-2 configuration eliminates the need for a reflective license plate altogether and simplifies the geometric computations.

In the IRIS-2 configuration, the algorithm for calculating the distance and yaw of the preceding vehicle can be developed using the similarly straightforward geometric method depicted in Figure 3. We define the coordinate with the origin (0,0) at the point S_2, which corresponds to the focal point of the right CCD sensor. L_{ll} denotes the line from the left reflective taillight on the back of the preceding vehicle to the focal point of the left CCD sensor installed on the following vehicle. The other three lines L_{lr}, L_{rl}, and L_{rr}, have the same geometric meaning as L_{ll}. Now it is straightforward to obtain the analytic expressions of these lines:

$$\begin{aligned} L_{ll}&: y - d = \frac{-u_{ll}}{f} x \\ L_{lr}&: y - d = \frac{-u_{lr}}{f} x \\ L_{rl}&: y = \frac{u_{rl}}{f} x \\ L_{rr}&: y = \frac{-u_{rr}}{f} x. \end{aligned}$$

The intersection point L of L_{1l} and L_{2l}, which stands for the position of the left taillight of the preceding vehicle, is computed as

$$\begin{aligned} x_r &= \frac{f \cdot d}{u_{lr} - u_{rr}} \\ y_r &= \frac{d \cdot u_{rr}}{u_{lr} - u_{rr}}. \end{aligned}$$

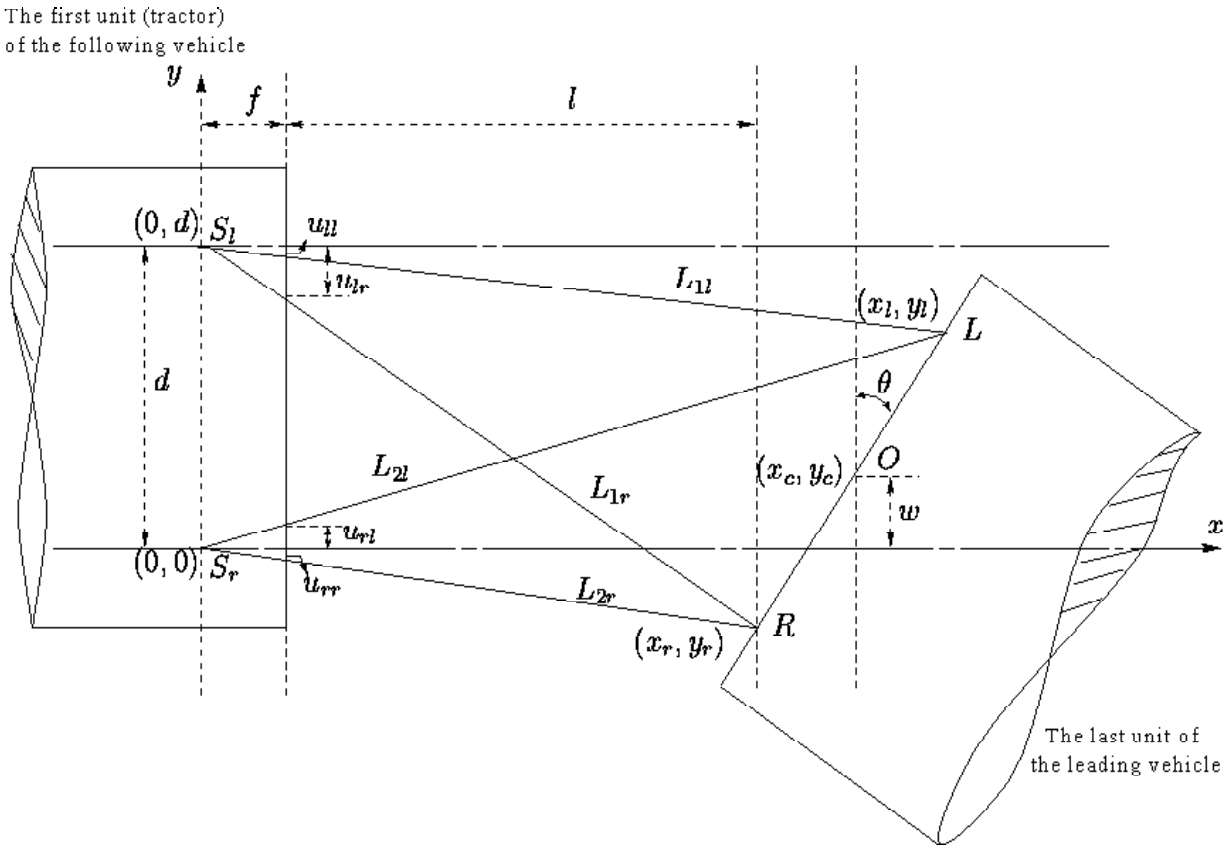

Figure 3

Geometry for extracting lateral and longitudinal information from IRIS-2.

Similarly, the intersecting point R of L_{1r} and L_{2r}, the position of right taillight of the preceding vehicle, is

$$x_l = \frac{f \cdot d}{u_{ll} + u_{rl}}$$
$$y_l = \frac{u_{rl} \cdot d}{u_{ll} + u_{rl}}.$$

From these equations, the longitudinal and lateral distances l and w and the yaw θ are computed as

$$l = \min\left(x_r, x_l\right) - f$$
$$w = \frac{y_l + y_r}{2}.$$
$$\theta = \tan^{-1}\left(\frac{x_l - x_r}{y_l - y_r}\right).$$

From the above procedure for computing the longitudinal and lateral distance and the yaw of the preceding vehicle, we see that IRIS-2 does not require knowledge of the distance between the two taillights of that vehicle, while that knowledge is required in the IRIS-1 configuration. Thus, a vehicle equipped with the IRIS-2 sensor can accurately obtain all the necessary information for longitudinal and lateral control for any type of preceding vehicle, as long as that vehicle has some retroreflective surface as required by law. This means that the IRIS-2 sensor can be used to detect and track not only cars and trucks, but also motorcycles and even cars with one broken taillight.

Close following. An important issue is the ability of IRIS to track vehicles at very small distances. To achieve this, it must be able to see both taillights of the preceding vehicle; even with its 80-degree field of view, it may lose one of the taillights at distances smaller than 2m (6ft). However, if two receivers are used, as in the IRIS-2 configuration, then the orientation of one of them can be appropriately adjusted to guarantee tracking at distances as low as 0.6m (2ft).

Reflective surfaces. Due to its operating principle of image subtraction, IRIS cannot detect any objects that do not have reflective surfaces. On the other hand, it can very easily and reliably detect any reflective object, such as vehicle taillights, street signs, overhead signs, and even reflective lane markers.

Adverse weather. Like any other sensor operating in visible or near-visible wavelengths, IRIS is not able to operate reliably in very heavy rain, fog, or snow. However, because it relies on the subtraction of two images and not on time-of-flight measurements, the operating weather threshold of IRIS is significantly higher than that of other infrared or vision-based sensors.

Multiple targets. Because each vehicle in the wide field of view of an IRIS sensor appears as a separate triangular pattern, it is very easy to track multiple targets without being confused by their separate returns. To achieve this, however, a fast and reliable cluster classification algorithm is required, which can identify the triangular patterns so they can be tracked.

Lane entry and departure. Due to the high reliability of its measurements and the robustness of its operating principle, the IRIS sensor is virtually immune to false returns. Hence, when another vehicle enters the lane in front or when the vehicle that was being tracked leaves that lane, the IRIS sensor immediately knows about it due to its ability to track lateral motion within its wide field of view.

1.4. Potential Uses

Using the above list of advantages and disadvantages, it is straightforward to conclude that IRIS will be useful in the following ITS applications:

Collision warning and avoidance. Clearly, when IRIS is used as a stand-alone sensor for collision warning, it can only warn the driver about impending collisions with other vehicles whose reflective surfaces are clearly visible. It will not be able to detect nonreflective obstacles such as pedestrians, trees, and highway barriers. However, its most important advantage over radar-based systems (which can detect nonreflective objects and can operate under more adverse weather conditions) is that, unlike IRIS, radar suffers from shadow returns that can set off

frequent false alarms and prompt the driver to turn off or ignore the system. To avoid that, such systems usually set a high return threshold above which an alarm will be triggered. Consequently, they often misjudge low-level returns from small or low-profile vehicles and motorcycles as false, and do not activate their alarm until these targets are much too close for effective action to be taken.

The IRIS sensor can significantly mitigate this problem when used as a **very-low-cost add-on** to such systems. By combining the radar returns with the virtually foolproof detection of IRIS, such a system would eliminate the threshold problem for vehicles and other objects with reflective surfaces in most weather conditions, and it would still be able to recognize nonreflective objects and operate in severe weather, albeit in somewhat degraded mode.

Adaptive cruise control. Since IRIS-1 can only measure changes in the distance from an arbitrary preceding vehicle, it can be used as a stand-alone ranging sensor for adaptive cruise control systems that use a **driver-initialized** spacing policy. In other words, the driver is responsible for first achieving the desired spacing and then activating the system; this way, IRIS-1 can be used to maintain that selected distance, even though it does not know its exact value.

It is well known that constant spacing results in less than ideal performance for adaptive cruise control, and that time-dependent spacing yields better response. This spacing policy, usually called **time headway**, expresses the desired spacing from the preceding vehicle in seconds and compares it to the current spacing, which is computed as the absolute distance divided by the current speed of the host vehicle. To implement this spacing policy, the controller needs to know the current speed of its own host vehicle (always available) as well as the absolute distance to the vehicle in front, which the IRIS-1 sensor cannot provide.

Nevertheless, the above principle of driver initialization applies to time headway as well. In particular, it can be shown that the time headway can be computed up to an unknown constant from the apparent separation of the taillights and the speed of the host vehicle. Once the driver achieves the desired headway and activates the system, it is possible to maintain this headway by controlling the apparent separation in response to changes in the vehicle speed; the unknown constant is not required for this process.

On the other hand, driver initialization is not as user-friendly as the ability to “dial up” a desired distance or headway. Consider, for example, the case where the preceding vehicle leaves the lane and a new vehicle must be tracked: in an IRIS-1 system, the driver has to repeat the initialization procedure each time. Clearly, it would be preferable to use an IRIS-2 system, which can close in on the new target without any additional input from the driver. Furthermore, IRIS-2 can be used to implement mixed spacing policies, which combine constant spacing with an added time headway.

It is worth noting that the capabilities of IRIS-1 can be augmented to obtain accurate absolute distance measurements and eliminate the need for driver initialization:

1. Fixed reflective patterns: One can require that all vehicles to be tracked are equipped with three reflective patches (like small bumper stickers) in a fixed-distance pattern. This low-cost solution is easily implementable in fleet settings (truck fleet operators, railroad cars), and it offers the additional benefit of being able to distinguish fleet vehicles from other vehicles on the road through the use of a distinctive reflective pattern.

2. Additional sensors: If the vehicle already has another sensor installed (microwave radar, laser radar, vision) which has the capability of measuring absolute distance, then IRIS-1 may be used as a very low-cost add-on for enhancing reliability and improving performance.

Vehicle following/electronic towbar. In the second stage of vehicle automation, when automated steering is added to automated throttle and brake operation, the control objective becomes more ambitious: Instead of merely alerting the driver to the presence of a vehicle in front or maintaining a desired spacing from it, one is now interested in actually following this vehicle through turns. Since IRIS can easily measure the lateral motion of the preceding vehicle, it can be used as a stand-alone sensor for automated steering. However, in that case the driver has to be certain that the vehicle in front can be trusted to follow the desired path. Therefore, this use of IRIS is very well suited for an electronic towbar scenario, where a commercial truck fleet operator forms an electronic train of several trucks, with the first one driven manually by a human driver and the others following in fully automated driverless mode. In a general highway setting, however, it may be preferable to use a system that follows the vehicle in front only as long as that vehicle remains in the current lane. A lane detection capability is then needed; this is discussed in the next paragraph.

Lane departure warning and road curvature prediction. The ability to measure the position of the vehicle in the current lane, as well as the curvature of the lane and the road ahead, is critical for applications such as lane departure warning and avoidance. Again, since IRIS can only detect reflective surfaces, it cannot be used as a stand-alone sensor for such an application in a general road setting. However, most roads and nearly all highways have some reflective surfaces to aid drivers in predicting road curvature at night. These surfaces include reflective paint for lane marking and roadside posts in Europe, and even roadside and overhead signs. Since IRIS can detect not only the presence but also the shape and the relative velocity of such surfaces, it can be used as a sensor for road curvature. This potential is further enhanced by the ability of IRIS to reliably monitor the movement of several vehicles in its wide field of view, even of vehicles that are several hundred meters away; after all, tracking the motion of preceding vehicles is commonly used by human drivers as an aid for predicting road curvature.

As is evident from the above list, IRIS-2 has significant potential as a low-cost stand-alone sensor. As for IRIS-1, its most important use may prove to be as a very low-cost add-on to be combined with another sensor such as radar or machine vision. Such combinations will result in enhanced reliability and a larger region of operation, and will be usable in a wide range of ITS applications.

2. PROTOTYPE SENSOR HARDWARE

The main accomplishment of this project was the construction of a working prototype of the IRIS-1 sensor. Figure 4 shows this prototype mounted on a wheeled cart; it consists of

- a CCD image sensor, a commercial infrared laser illuminator mounted on a two-axis tilt mount so that the beam direction can be adjusted to coincide with the field of view of the image sensor, a relay lens that sits between the front camera lens and the mask emulator (see below), and a bandpass optical filter that was placed on the front camera lens in order to filter out all the incoming light except that which has the same wavelength as the infrared illuminator;

- a black plastic shroud that encloses the entire optical setup, in order to partially shield the light from the sides and to ensure that the operation of the system is eye-safe, by preventing anyone from placing their eye directly on the output of the laser; and

- a 120-MHz Pentium laptop PC that controls the hardware, interfaces with the CCD camera to acquire the image data, and processes this data to identify the taillights and determine the current state of the target vehicle.

This prototype was first used in static and dynamic experiments indoors and outdoors, with the sensor mounted on the wheeled cart shown in Figure 4, which was being pushed manually. For the final project demonstration, this sensor was removed from the cart, mounted inside a vehicle, and operated to compute the relative distance to another moving vehicle in real time.

2.1. Mask Emulator

The IRIS sensor concept is based on a masked CCD array that allows the acquisition of two images in rapid succession. In production devices, the desired mask can be implemented during device processing directly on the CCD chip. Such a process can be easily implemented using existing technology, since it involves only a subset of the steps required for the manufacturing of color CCD chips. However, it would require the modification of existing production lines, and therefore the cost of producing just a few prototypes would be prohibitively high. To circumvent this problem, we built a prototype image sensor in which the masking is realized using a two-step optical system. The scene is first imaged onto a mask, and then the masked image is re-imaged onto a standard commercial CCD array built by Pixel in Oakland, California.

Our CCD array is the TC 245 made by Texas Instruments. The active area of this chip has 242 rows of 756 pixels. Each pixel is 19.7 μm high and 8.5 μm wide. This is an image-transfer device in which the active area is shifted rapidly into a storage region before read-out. We chose this chip as the basis of the prototype IRIS sensor for the following reasons:

1. The large number of pixels in the horizontal direction (756) allows high resolution of the horizontal positions of the target vehicle retroreflectors, as is required for accurate range determination.

Figure 4

Working prototype of IRIS-1 sensor.

2. The ability to transfer the image rapidly into a storage area before read-out allows images with bright regions to be acquired at short exposure times without excessive smearing, even without the use of an external shutter: Since the pixels continue to accumulate charge as long as they are exposed to incident light, if the image is read while still being exposed, the brightest regions of the scene can leave an increasingly bright vertical track across the recorded image. In an image-transfer chip the image is rapidly shifted into a masked storage area and then read out, so that the time the image is exposed during the read-out process is minimized.

3. A variety of interface options exist for this chip, since it is widely used in both commercial and home-built imaging systems. Thus, computer interface electronics are readily available and several examples of source code can be used as patterns for the custom control software required for our application.

The optical elements in the prototype imager include

- a bandpass filter (10 nm bandwidth centered at 800 nm),
- an imaging lens (f = 6 mm, 1:1.2),

- the mask (a 300 line-per-inch Ronchi ruling), and
- a relay lens (f = 50 mm, d = 25 mm).

These elements are mounted in a structure based on ThorLabs components to provide the rigidity required in an automotive environment.

The sensor setup is shown schematically in Figure 5. The combination of the Ronchi ruling and relay lens results in the masking of pairs of rows on the CCD array, rather than single rows (the best solution available with off-the-shelf components). The loss of vertical resolution resulting from this masking pattern is not critical, since it is primarily the horizontal resolution that is used to compute the range. The choice of relay magnification was determined by the availability of suitable commercial components (lenses and Ronchi ruling). The resulting masking scheme is shown in Figure 6.

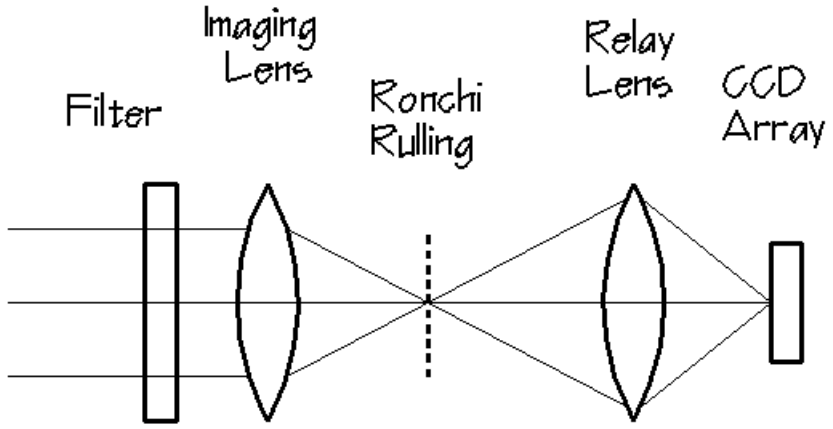

Figure 5:

The optical setup of our sensor prototype.

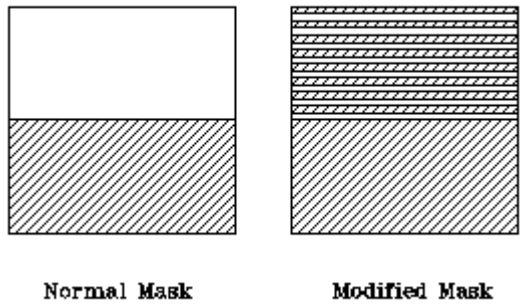

Figure 6:

Schematic showing the original mask configuration of our CCD chip and the effective masking scheme achieved by the setup of Figure 5.

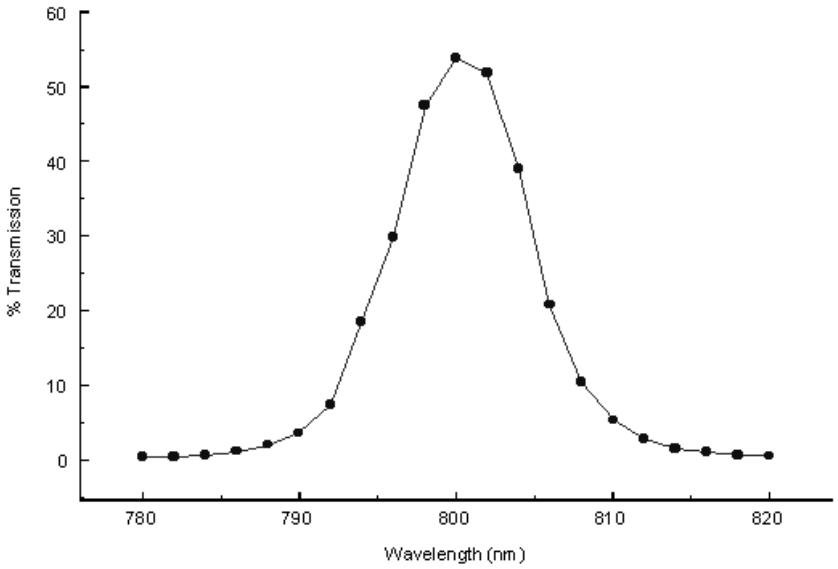

Figure 7:

The transmission characteristic curve of our bandpass filter.

The focus of the relay lens and the alignment of the Ronchi ruling lines with the rows of the CCD array are crucial to the proper function of the image sensor. We emphasize again that this is the portion of the prototype that performs the masking function, and it will not be required when a pre-fabricated masked CCD array chip is available. The mechanical tolerances for this portion of the image sensor are extremely tight. The lines of the Ronchi ruling must be imaged at the correct magnification, they must be translated vertically to match the 19.7-μm-high pixels, and the rotation of the lines relative to the rows of CCD pixels must be correct to a few μm over a lateral distance of about 6.4 mm. We mounted the Ronchi grating on a two-axis tilt mount to accomplish this alignment. One axis of rotation allows the lines of the grating to be tilted with respect to the pixel rows. The other rotation axis tilts the grating with respect to the optical axis, thus allowing precise shifts of the vertical registration of the grating with respect to the pixel rows. This mechanism has been constructed and the grating has been aligned with the CCD array. In order to filter out the lights from some strong light sources, an optical band-pass filter is incorporated in the image sensor prototype. The transmission of this filter as a function of wavelength is shown in Figure 7.

2.2. Laser Illuminator

The IRIS illuminator utilizes a commercial laser diode at a wavelength of approximately 800 nm, which is matched to the transmission of the optical bandpass filter in the image sensor. The total output power of the laser is 0.5 W. This power is emitted into a highly asymmetrical beam, which appears as a stripe in the far field. (The half-angles of the horizontal and vertical directions are approximately 40 degrees and 5 degrees respectively.) The illuminator uses two lenses to produce an output beam that illuminates a field of view matched to the field of view

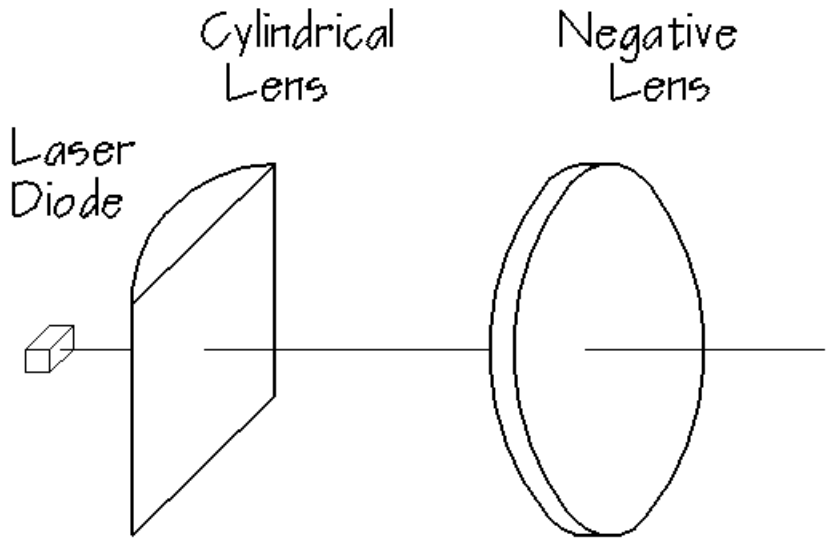

Figure 8:

The illuminator setup of our experimental prototype.

of the image sensor. A diagram of the illuminator optics is shown in Figure 8. A cylindrical lens is used to compress the beam divergence in the horizontal direction so that the vertical divergence is about one-third that in the horizontal direction, rather than one-eighth. Then a negative lens is used to uniformly increase the divergence of the entire beam pattern. This optical assembly is mounted on a two-axis tilt mount so that the beam direction can be adjusted to coincide with the field of view of the image sensor.

A commercial laser diode driver package supplies power to the laser. The driver electronics include a control line that allows the laptop computer to turn the illuminator on and off.

3. PROTOTYPE SENSOR SOFTWARE

3.1. Driver and Image Acquisition Modules

In order to read and process the images that are produced within the CCD array, a production version of the IRIS sensor would use an integrated microprocessor connected to the CCD chip in a small package. For the purpose of development, however, we also needed to be able to view and record these images. Thus, in our prototype system we connected the array interface electronics to a parallel port of a laptop computer and used a software module to control the illuminator and receiver through this connection. Since the way we use the CCD camera is unique, it was necessary to write our own software to provide this control function. The first version of this program was a DOS application and was used only long enough to get the signal sequencing correct. At that point it became apparent that a Windows application would be able to display images with less programming effort, more quickly, and with the ability to be more easily manipulated. Therefore, the DOS application code was ported to a 16-bit Windows application called Win245. The "245" refers to the Texas Instruments TC 245 chip used in the image sensor.

The Win245 program is a 16-bit Windows application with a graphical user interface that can not only communicate with a CCD camera attached to the LPT1 parallel port, but can also display and manipulate the images produced by the camera. The program takes care of all the timing and sequencing of the control signals that must be sent to the camera in order to acquire images. Typically, the program takes one picture at a time, but it can take several images in sequence if desired. For each sequence of images, the user can specify the exposure time, whether or not the illuminator should be used, whether or not the expose-shift-expose procedure should be used, and whether or not the images should be subtracted automatically.

The program refers to the use of the expose-shift-expose procedure as "using the grating." In this mode, each image is created by first emptying all the exposed pixels of the CCD camera and then allowing them to gather light for the specified exposure time. The pixels are then shifted down by a user-specified number of rows (this number is equal to 2 in our current experimental setup) and then exposed a second time. With the grating in place, this mode allows one image to be taken first and then stored temporarily in the rows shielded from exposure by the grating while another image is taken in quick succession. If the user specifies that the illuminator is to be used as well, the first image will be taken with the illuminator on while the second image will have the illuminator off. If image subtraction is used, the second image will be subtracted from the first before proceeding. Only this subtraction result is stored. If subtraction is not used, then both images are stored.

When the expose-shift-expose procedure is not used, the pixels are simply emptied and exposed for the specified exposure time. In this mode, use of the illuminator just entails turning the illuminator on and off from one image of a sequence to the next. The first image always has the illuminator on. Image subtraction in this mode means that two exposures are taken for each image of the sequence, with the first image being completely read from the camera before the second is taken. The second image is then read and subtracted from the first.

Once an image sequence has been completely generated, the program displays each image as a separate window on the screen. Each of these image windows comes complete with its own scroll bars and cross hairs that follow the mouse so that real-time horizontal and vertical cross sections can be viewed in the horizontal and vertical cross section display boxes. The horizontal cross section display is just below the image display while the vertical cross section display is just to the right of the image display. Examples of such images are shown later on in Figures 13 and 15.

The user has several options available for working with these resulting image sequences. An image calculator is available that can add, subtract, multiply, or divide two images on a pixel-by-pixel basis, or a single image can have one of these operations performed on it using each pixel and a user-provided constant as the arguments. This allows the user to perform, for instance, an image subtraction after the images have been taken. Each image can be individually zoomed in once, providing a 5x horizontal and 2.5x vertical magnification. Images can also be "viewed as text" in which the visual representation of the image is replaced by rows and columns of numbers that represent the intensity seen by each individual pixel of the camera.

There are two options available for saving image sequences. The first is the "Save As" option that saves all images currently displayed in a binary format. The second is the "Save As Text" option that saves all images currently displayed in a text format that can be read by most word processing or spreadsheet programs. In both cases, the program stores each image in its own file whose name is created using the first five letters of the user-provided file name and three numbers starting with 000. For example, if the provided file name is "picture" and there are five images currently displayed, then the first image will be saved as "pictu000.245" and the fifth as "pictu004.245." The program is only capable of opening the binary formatted files, so everything should always be saved in at least this format. When a file is selected for opening, the program will search the directory for every file with the same first five letters and open all of them in order. This makes keeping image sequences together much easier.

As we mentioned above, our CCD camera has 242 rows of 756 pixels. The bottom 121 rows are covered by the protective metallic mask of Figure 6, so that leaves the top 121 rows as the active area where the CCD image is recorded, for a total of $121 \times 756 = 91{,}476$ active pixels. When the CCD image is transferred to the laptop computer, each of these pixels is read out in series, after first passing through a 12-bit analog-to-digital (A/D) converter. This results in a total of $91{,}476 \times 12 = 1{,}097{,}712$ bits for each downloaded image. With the additional communication overhead, the total number of bits that have to be transferred through the parallel port of the laptop for each image is about 1.2 million. Since the transmission speed of the parallel port is approximately 400 kb/sec, we see that it takes 3 sec to download each image. Therefore, with the camera running in full resolution mode, the refresh rate of our prototype is 1/3 Hz. To increase the rate, we used a technique called "binning" to reduce the number of pixels that had to be read from the camera. Due to the physical properties of CCD cameras, it is possible to merge the contents of one pixel with those of another during the readout process. This merging of pixels is called "binning" and allows the resolution and dynamic range of the camera to be reduced in order to increase the readout speed. For the IRIS-1 prototype, every three horizontal pixels along with every two vertical pixels were binned, resulting in a six-fold reduction in resolution and an approximate six-fold increase in update rate. This increased the update rate from 1/3 Hz for the full-resolution mode to 2 Hz for the low-resolution binned mode, which was used for our in-vehicle moving experiments.

Furthermore, the image acquisition module included a subroutine that computes the histogram of the pixel intensities during the image transfer from the CCD camera to the laptop. This histogram was then passed along with the image to the object detection module, where it was used to determine the threshold value for noise filtering; this process is described in detail in the next section.

3.2. Object Detection Module

The driver module operates the CCD camera to record two successive images closely spaced in time, the first one with the laser on and the second one with the laser off, and store them both in the bottom part of the chip where they are protected from further exposure to light by the metallic mask. Then, the image acquisition module transfers these two images to the computer memory, subtracts the second one from the first one, and creates a histogram of the pixel intensities of the resulting subtracted image. This data is then passed to the object detection module, which processes the image data to detect all recorded objects with strong returns and determine their center positions. The three main steps in this are: (i) suppressing the noise, (ii) distinguishing between separate targets, and (iii) determining the center location and intensity of each target.

Noise Suppression. When a return becomes too bright, a phenomenon called "blooming" occurs in the CCD sensor array, which can be described as follows: The individual pixels act like charge-holding cups. The charge (number of electrons) in each "cup" is proportional to how bright the return is at that pixel. When the return at a CCD pixel gets too bright, the charge fills to the top and, just like a cup, overflows into neighboring "cups." Such overflows may occur in only one of the images (laser on or laser off). In these cases, subtraction of the first (laser off) image from the second (laser on) image cannot eliminate the overflow. This effect accounts for noise in the resulting subtracted image. The suppression of such noise is imperative to the performance of the object recognition algorithm that follows. Since the overflows are in general weaker than real returns from retroreflective surfaces, they can be easily suppressed by **thresholding**.

For our first experiments, we selected a threshold manually, by looking at each image to determine where to place the threshold. However, for realistic in-vehicle experiments all processing has to be done in real time without human intervention. This means that the threshold value for each image must be determined automatically by the software as well. This selection was performed via a well-known image processing technique that uses the histogram of the distribution of gray levels in the image; this histogram is created by the image acquisition module during the image transfer process. Because background noise (return from surfaces other than the retroreflectors) is much lower in intensity than the return from the taillights, the histogram will generally have two major hills. One starts from the beginning of the histogram (at a very low gray level); the next starts where the intensity of the taillights lies. An example of the type of histograms we obtain is shown in Figure 9. The threshold should therefore be set around the point where the second hill starts. Effective thresholds can be determined instantaneously with real-time-updated histograms, provided an efficient algorithm can be developed for locating the beginning of the second hill.

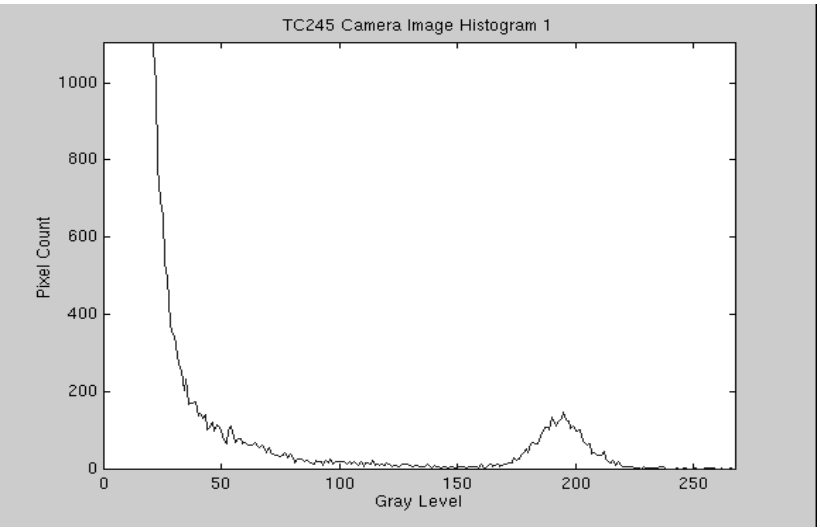

Figure 9

Typical histogram before smoothing.

The algorithm we developed looks for one large drop followed by a large rise to determine the threshold. Even though the hills we are looking for are much larger compared to other bumps in the histogram, we cannot assume a continuous drop or rise where the hills occur. In fact, almost all histograms exhibit jagged rises and drops like the ones seen above. To eliminate this problem, the histogram is first passed through a low-pass-filter. The scheme used here replaces the pixel count at a particular gray level with the average pixel count of several neighboring gray levels. The resultant histogram is almost free of inconsistent drops or rises. After this filtering process, the histogram of Figure 9 is much smoother, as seen in Figure 10.

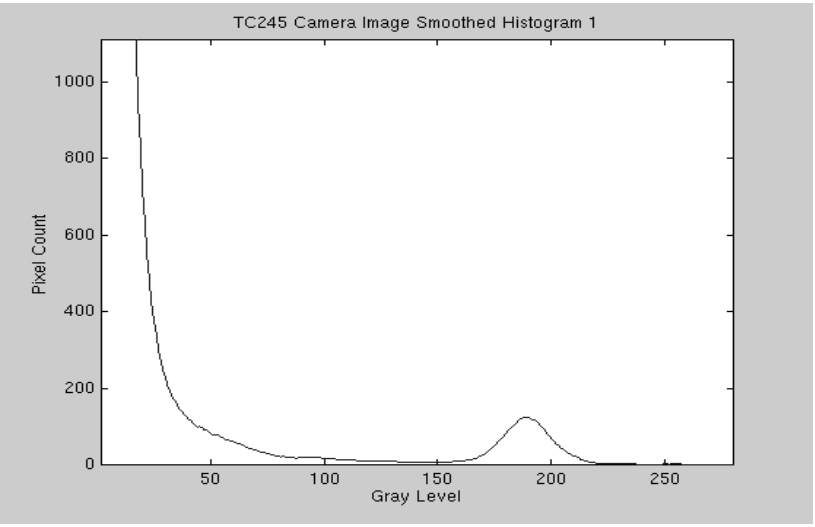

Figure 10:

The histogram of Figure 9 after smoothing.

Smooth histograms allow us to count the number of drops and rises to determine where the threshold lies. The following is a detailed algorithm description, where *navg* is the number of pixel counts to average over, *ndrop* is the number *drop_count* has to reach to declare a large drop found, and *nrise* is the number *rise_count* has to reach to declare a large rise found. A flowchart of the algorithm is given in Figure 11.

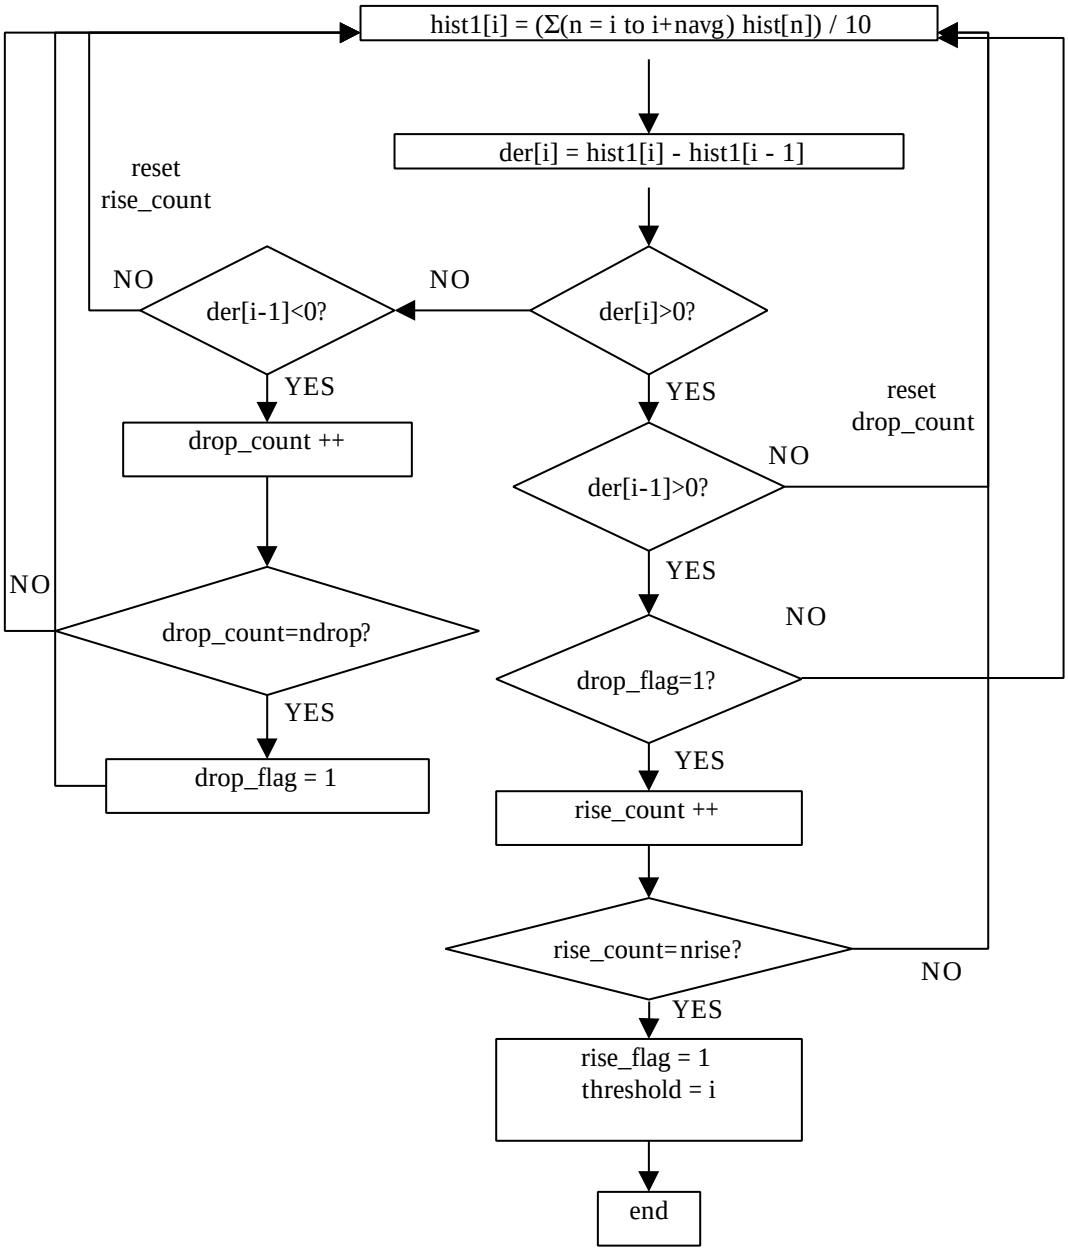

Figure 11

Flowchart of threshold algorithm.

1. Replace current pixel count by the average of the current and the next *navg* pixel counts.
2. Subtract the previous pixel count from the current one (calculate a discrete derivative).
3. If the current and the previous derivatives are both less than zero (dropping), then increment *drop_count*.
4. If the current derivative is dropping, and the previous derivative was rising, then reset *rise_count* to zero.
5. When *ndrop* drops are reached, declare a large drop found.
6. If the current and previous derivatives are both greater than zero (rising), increment *rise_count* only if a drop has been found.
7. When *nrise* rises are found, declare a large rise found, and the threshold to be the current gray level.
8. Calculate the next histogram value by subtracting off the pixel count from the lowest gray level being averaged, and adding one pixel count beyond the highest gray level added last time, then dividing by *navg* again.
9. Repeat steps 2 – 8 until a rise is found or the end of histogram has been reached.

Pattern Separation. The problem of determining the center points of all objects present in the image has two parts: (a) recognizing individual objects, and (b) finding their respective center positions. Since other objects such as license plates, road signs, and taillights from other cars are also sufficient retroreflectors at the illuminator's wavelength, even the most basic model of just one car may have three distinct objects after thresholding.

Utilizing some functions in Matlab, an iterative algorithm was developed to separate all disjoint objects in a given thresholded image. The *find* function in Matlab scans a given matrix one column at a time, returning the positions and values of the non-zero elements. *[row, col, vlu] = find(A)* returns, for all nonzero elements in matrix A, the row numbers in vector *row*, column numbers in vector *col*, and values in vector *vlu*. This function provides, with reasonable processing, all the necessary information for finding the center positions of all disjoint nonzero regions in an image.

An object is defined as one with no complete vertical or horizontal breaks. The first step is to recognize which information in *row*, *col*, and *vlu* belong to one object. Since there is an unknown number of disjoint objects in the image, an iterative algorithm was selected. The following is a detailed algorithm description. Figure 12 shows the flowchart of the pattern separation algorithm.

1. Given an image, find the first completely zero row and/or column (called breaks from now on).
2. Cut the given image vertically and horizontally (into four pieces) if both vertical and horizontal breaks exist. If only horizontal (vertical) break exists, cut horizontally (vertically). If no breaks (neither horizontal nor vertical) are found, declare the given image to be one object.
3. Find the positions (*srow*, *scol*) and values (*svlu*) of the non-zero pixels in the resultant sub-image(s). Add row and column offsets to *srow* and *scol*, respectively (row and column offsets for the first iteration are zero).

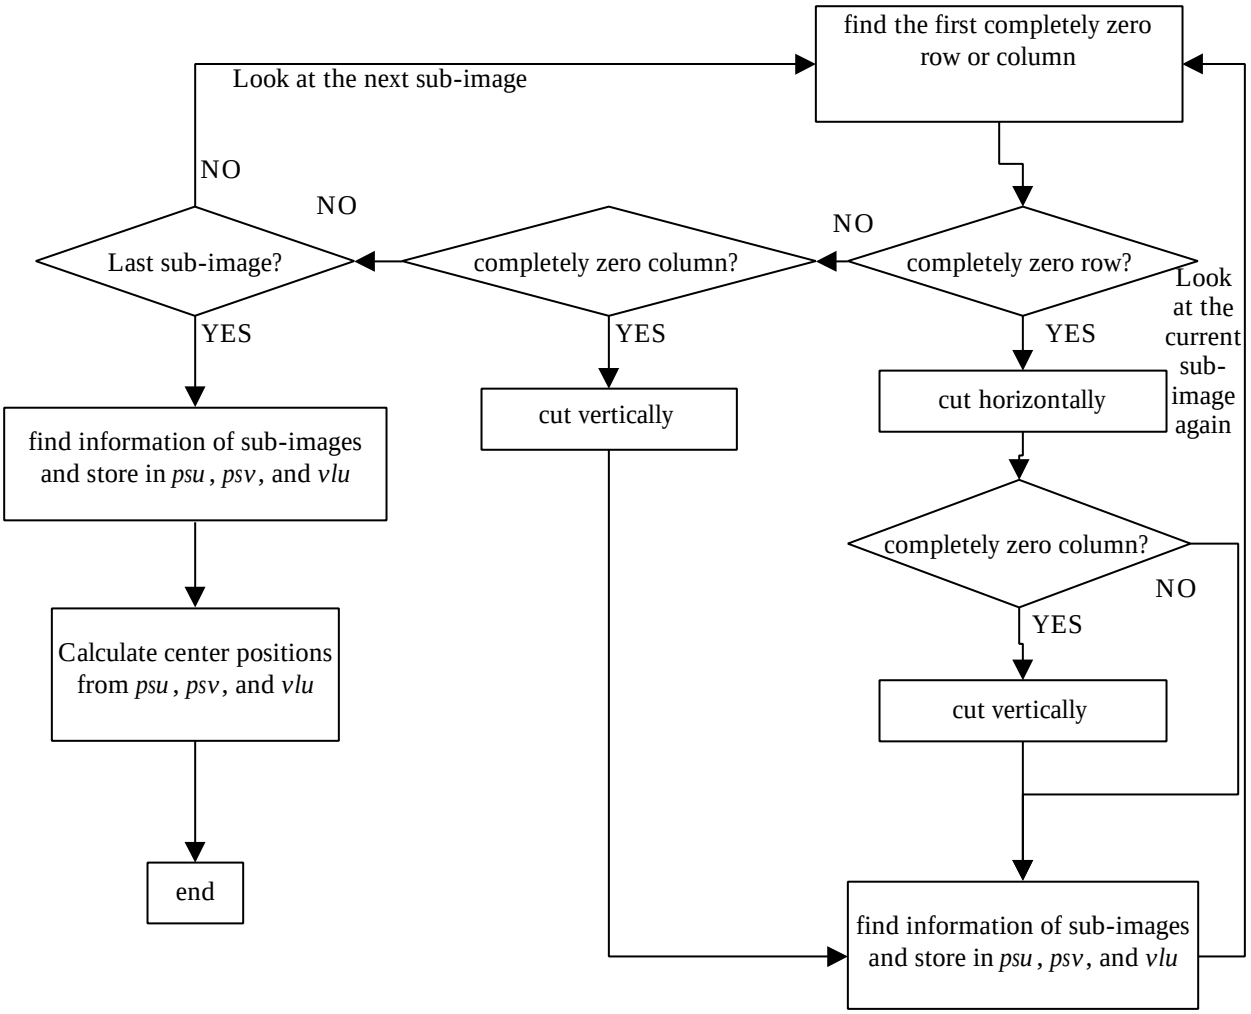

Figure 12

Flowchart of pattern separation algorithm.

4. Put *srow*, *scol*, and *svlu* into the final information matrices *row*, *col*, and *vlu*. *Row*, *col*, and *vlu* are defined as matrices, where the corresponding columns in the matrices contain the information of the non-zero pixels in one sub-image. Here,

   L = number of columns in row (col, or vlu);
   row = [row(:, 1:i-1) srow row(:, i+1:L)];
   col = [col(:, 1:i-1) scol col(:, i+1:L)]; and
   vlu = [vlu(:, 1:i-1) svlu vlu(:, i+1:L)].

   Effectively, we replace the information of the *i*-th sub-image with the information of its sub-images.
5. If the previous sub-image is already one object, then look at sub-image i+1; if not, look at the first resultant sub-image (sub-image i).
6. Determine the row and column offsets of the current sub-image (with respect to the original image), repeat steps 1 – 5 for the current sub-image.

7. Repeat steps 1 – 6 until all sub-images contain just one object (i.e., until each column of *row*, *col*, and *vlu* only contain information for one object).

This algorithm was designed to take into consideration that the shape of the taillights can vary from one vehicle to another, and that the apparent size of the taillights on the CCD image can change significantly with the distance between the host vehicle and the preceding vehicle. Since the algorithm only finds pixels that are not disjoint with respect to one another, it makes no assumption about the shape and size of the objects present in the given image; this makes it more robust with respect to such variations. However, if a small break (vertical or horizontal) is present in an object, the algorithm would recognize this object as two (or more) separate objects. The solution to this problem is to increase the number of completely zero rows (or columns) that can occur before declaring a break at that point.

After the row, column, and intensity information of the thresholded image is organized into matrices, the centers of the objects can then be found by performing a weighted average. For every column of the matrix *row*, and the corresponding column in *col* and *vlu*, the following is performed:

1. Every row and column number is weighted by the pixel intensity at that position.
2. The weighted row (and column) numbers are summed together.
3. The pixel intensities are also summed together.
4. The weighted row (and column) sum is divided by the intensity sum to obtain the final calculated center position.

The return from the taillights is usually brightest at the center of the taillights, then tapers off quickly. Using a weighted average to calculate center points thus presents significant advantages. Since the return is brightest at the centers, the center positions are weighted more heavily, making the calculation more accurate. Weighted average also makes no assumption on the size and shape of the objects.

For the IRIS-1 prototype implementation, the object detection algorithm described above was ported to our laptop computer and integrated with the other software modules. Its output, which was fed into the ranging module, consisted in the center positions and total pixel intensities of each detected object. As we will see, this is all the information that the ranging module needs in order to compute the distance from the preceding vehicle.

3.3. Ranging module

The main purpose of our prototype was to compute and continuously update the relative distance of the host vehicle from the target vehicle. In a complete product, this computed distance would be fed to another module that would issue a warning to the driver (in a collision warning system) or would activate the throttle or brake (in an adaptive cruise control system). For the purpose of our demonstration, however, this computed distance was displayed prominently on the screen of the laptop computer. The process of displaying the subtracted image and the computed distance is much more computationally intensive than the distance computation algorithm itself, but this was necessary for debugging and demonstration. The ranging module

can operate equally well in both the high-resolution-low-frequency and the low-resolution-high-frequency modes, which were described above. Since the image data transfer from the CCD to the memory of the laptop through the parallel port is by far the most time-consuming step in the entire process, the update rate is determined almost entirely by the resolution in which the camera operates; the time required for the range computation and display is negligible in comparison.

As mentioned above, the input to the ranging module consists in the center position and total pixel intensity of each detected object. All the retroreflective surfaces within the viewing range of the IRIS sensor show up as objects, including the taillights and license plate of the preceding vehicle and of the vehicles in neighboring lanes, traffic signs, lane markers, etc. For the purpose of this project, however, we wanted to test the ability of the sensor to track the vehicle in front, whose taillights always produce the two brightest returns on the CCD image. Therefore, our ranging module selected the two objects with the highest total pixel intensities (brightness of return), and assumed that they were the taillights of the vehicle in front. Then, using the simple trigonometric relations depicted in Figure 2, it computed the distance to the preceding vehicle using the known distance between the two taillights, which was 1.1 m.

A critical aspect of the ranging software module is the calibration of the triangulation parameters, which for this project include the focal length and the measurement bias. This calibration was performed off-line, using several images of the target vehicle at known distances away. Then we performed a least-squares fit to compute the best values for these two parameters. Of course, in a commercial product this calibration procedure will have to be automated.

4. EXPERIMENTAL RESULTS

4.1. Static and moving tests using a wheeled cart

For the first suite of moving tests, we mounted the IRIS prototype system on the wheeled cart shown in Figure 4 and took measurements outdoors by pushing the cart at various speeds while viewing the rear of an automobile. The images used an exposure time of 100 ms, that is, the CCD array pixels were cleared and exposed for 100 ms. The resulting pixel charges were then shifted down two rows and exposed for a second 100 ms period (the shift time of 100 μs is negligible compared to the exposure time). After the second exposure, the image charges were shifted to the storage portion of the chip, read out, and the subtracted image produced by the Win245 software.

Figure 13 shows the screen display of the Win245 program for a subtracted image taken while the cart was approximately 6.5 m behind the car and moving toward it at 3.2 m/s. A plot of the intensity vs. pixel position is shown in Figure 14 for the portion of this image containing the vehicle. Clearly, this experiment represents one of the worst possible scenarios, and we are not even sure that this could ever happen in an actual car; keep in mind that this cart is rolling on hard wheels without any suspension, and it is being pushed by a human who is running at 3.2 m/s. This mode of operation introduces a large amount of vertical vibration, whose effect is further compounded by the fairly long exposure time we had to use due to the limitations of our present setup. This causes the ambient scene in the two successive images to be quite different, thus resulting in imperfect cancellation of ambient noise after the image subtraction process. In a car driving down a highway with a regular suspension, an IRIS system using about one-quarter of the exposure time would never have to deal with anything this radical. Still, we see that the retroreflectors in the taillight covers are clearly identifiable above the considerable noise. The license plate is also visible, although its reflectivity is obviously lower than that of the retroreflectors, and its image is buried in the ambient noise. We used this data set as a benchmark for our ranging algorithms, to see whether our filtering techniques can detect the presence of the license plate and pick out its position from the noise.

Our conjecture that the poor (yet still usable) quality of the data in Figures 13 and 14 was caused by the fast motion of the cart, is supported by the data shown in Figures 15 and 16, which resulted from a similar test of the cart-mounted IRIS system when the cart was being pushed at the slower speed of 2.2 m/s. Note that the subtraction is much more effective in this image, due to the significantly smaller vertical vibration of the cart, which made the representation of the ambient scene in the two successive images much more similar, thus allowing it to be canceled out by the subtraction. The resulting signal-to-noise ratio of the reflective parts of the image vs. the non-reflective ones is very large, and we believe that this data set is much more typical of what we can expect in a real highway environment. Once again, the license plate is much less prominent than the taillights, although now it is very clearly identifiable. We believe that this difference is due to the fact that the reflectivity of the paint in the license plate is much lower at the infrared illuminator wavelength than that of the taillights. Both data sets clearly demonstrate that the image subtraction is effective even when the IRIS system is on a vibrating moving platform.

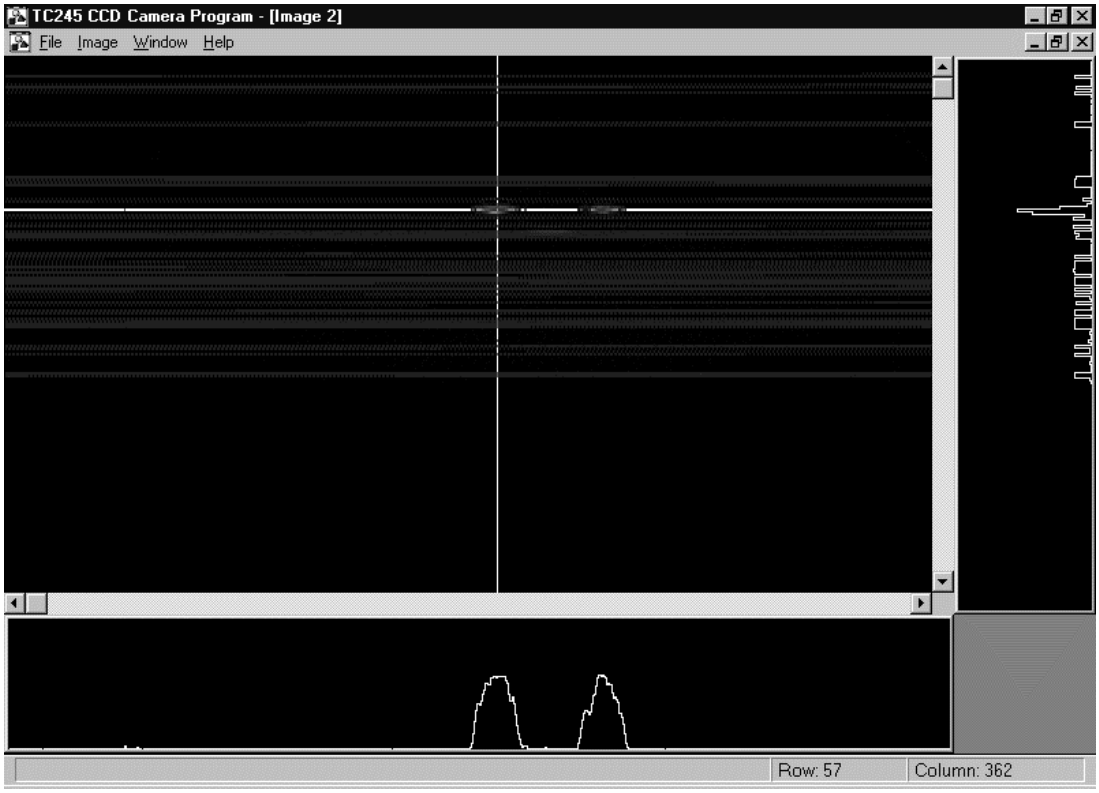

Figure 13

The Win245 image produced with the cart moving at 3.2 m/s.

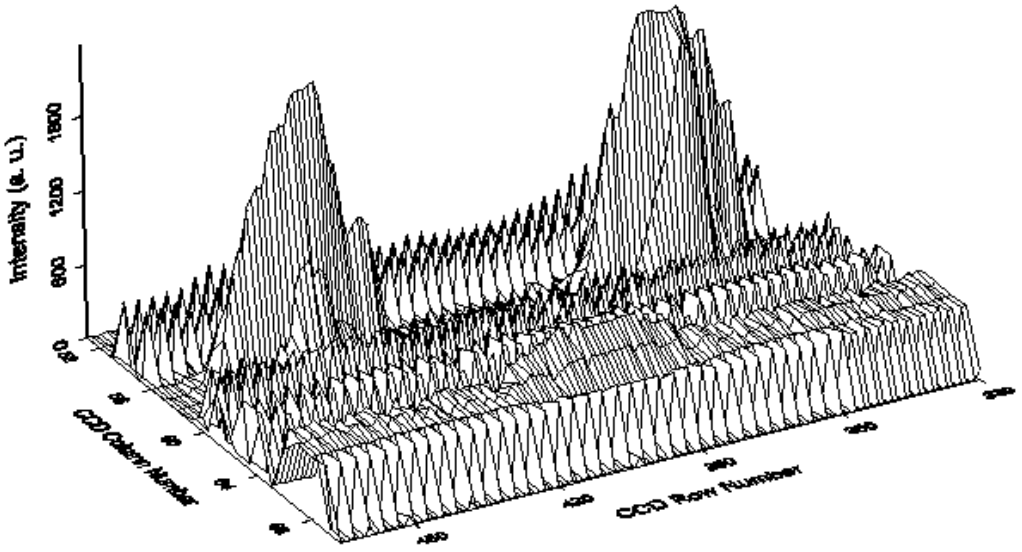

Figure 14

The intensity plot produced with the cart moving at 3.2 m/s.

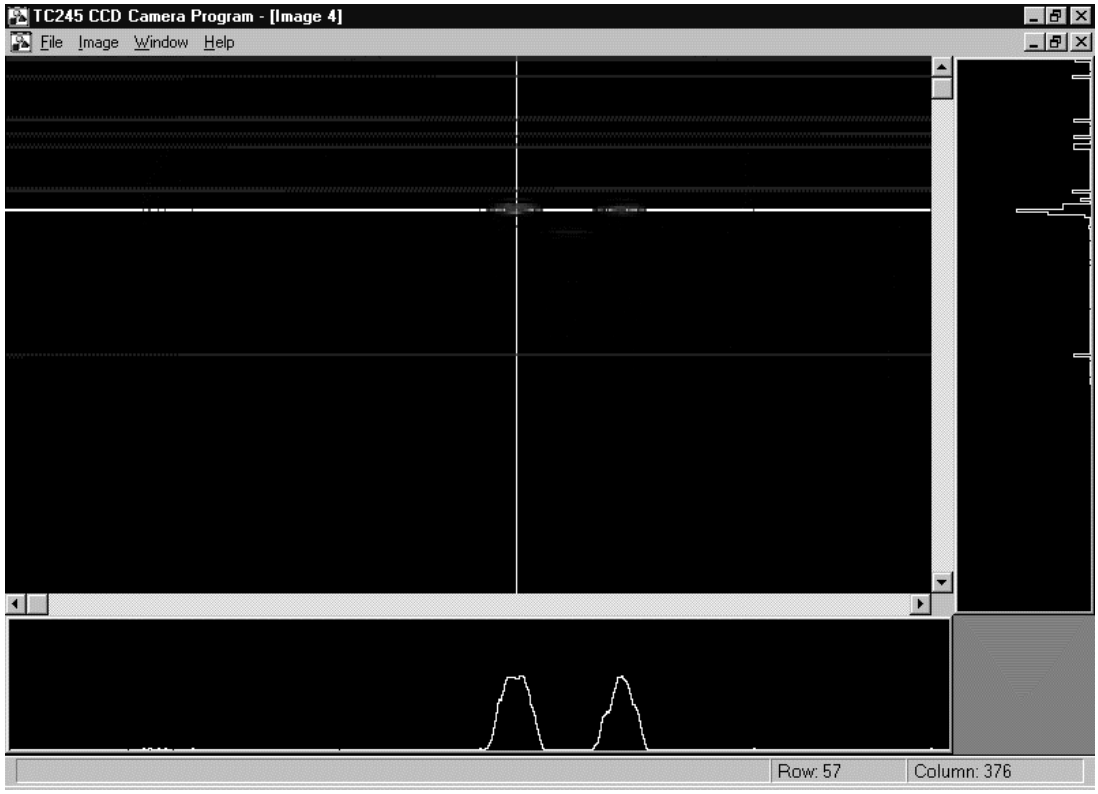

Figure 15

The Win245 image produced with the cart moving at 2.2 m/s.

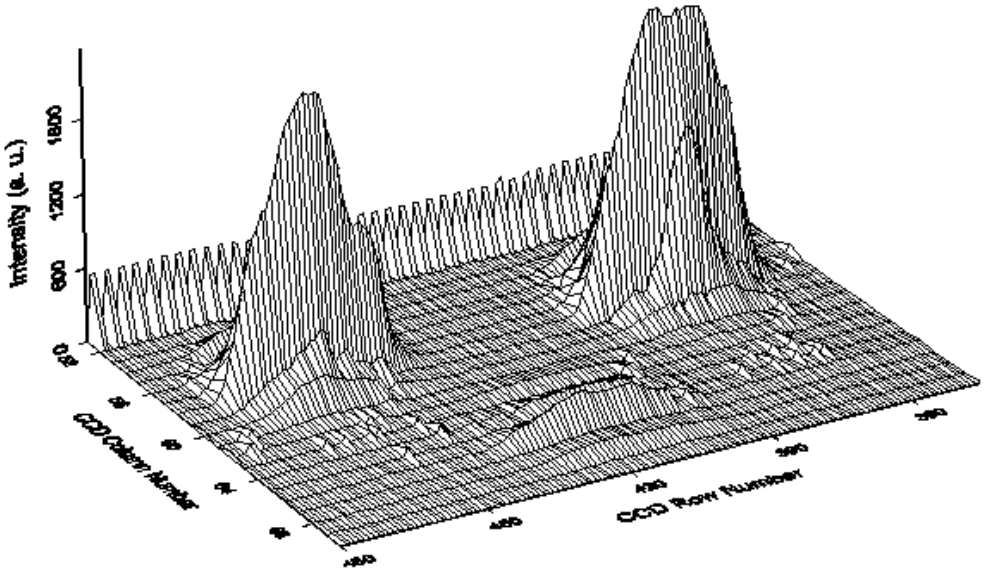

Figure 16

The intensity plot produced with the cart moving at 2.2 m/s.

To further test our conjecture, we performed an array of static tests in a laboratory environment, running the camera in the high-resolution (242×756) mode. The target was placed at distances varying from 5 ft to 60 ft, and the error of the computed range was consistently below 1%. This strengthened our confidence in the validity of our methodology, and naturally led to the next suite of tests, which were performed using two moving vehicles.

4.2. Tests using moving vehicles

On September 21, 1998, a demonstration was conducted on the UCLA campus in order to showcase the ranging capabilities of our prototype IRIS-1 sensor to the members of our Overview Technical Committee. The members of the committee were:

- Robert Diller, Vice President for Research of Interactive Voice Systems Inc.,
- Dr. Barry Dunbridge, Director for Electronics Technology in the Center for Automotive Technology of TRW Automotive, and
- Rudolf Lorenz, Director of the Vehicle Systems Technology Center of Daimler-Benz Research and Technology North America.

Due to the size of the prototype shown in Figure 4, it was impossible to mount it on the front side of any vehicle. Therefore, the whole setup was taken off the wheeled cart and placed in the rear cargo area of a Ford Explorer, whose cargo door was left open with the sensor pointing backwards. The prototype was powered from the electrical system of the car using a 12VDC-to-110VAC power inverter. The target vehicle was a Ford Escort, which was driven behind the Explorer and facing in the opposite direction, in order for the IRIS sensor to be able to track its taillight assembly. The suite of experiments included ranging with both vehicles stationary, with the Explorer stationary while the Escort was driven in reverse (closing in) and forward (moving away), with the Escort stationary while the Explorer was driven in reverse (closing in) and forward (moving away), and finally with both vehicles moving at the same time, at speeds up to 10 mph. The exposure time was set to 30 ms, and the camera was operated in the low-resolution (121×252) mode, which increased the update frequency to 2 Hz. At this frequency, the IRIS-1 prototype was able to accurately and reliably track the target vehicle at relative speeds of up to 10mph, with errors that remained smaller than 2% for distances ranging from about 6 ft to about 50 ft. The acquired CCD image and the corresponding computed range (in inches) were continuously displayed on the screen of the laptop computer, as depicted in Figure 17.

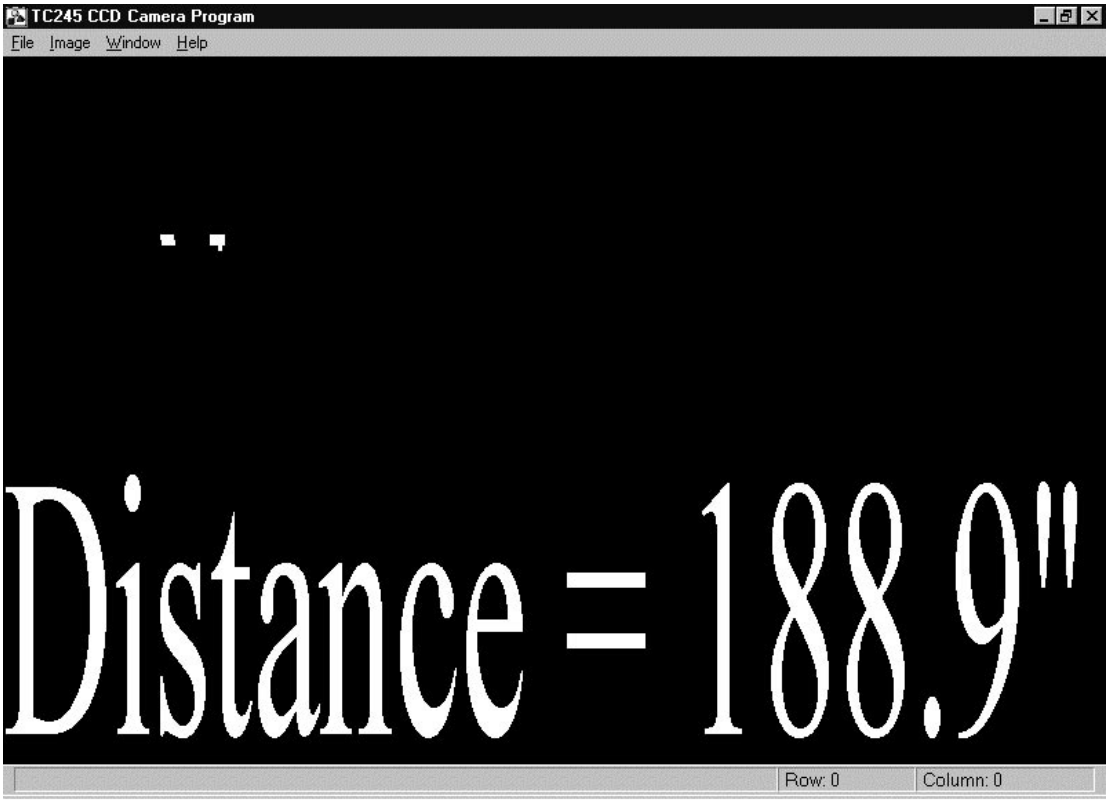

Figure 17

Display screen for IRIS-1 demonstration.

5. CONCLUSIONS

The ITS-61 project accomplished its main goal, which was to prove the feasibility of the IRIS concept. In fact, the performance of the prototype IRIS sensor performance exceeded our expectations even in a realistic in-vehicle environment. Of course, this success was only the beginning. Now that we know that the IRIS technology works, we are planning to pursue a broad array of research directions in order to realize the full potential of IRIS. Using some additional funding from TRW and Daimler-Benz, we are already working on resolving some of the problems we encountered during this project and implementing some of the many possible enhancements and extensions. This research is beyond the scope of the current project, but it will leverage its accomplishments, increase its payoff, and make this technology of even greater value to the ITS community. These optional tasks include:

- Installation of a new IR bandpass filter with improved rejection of out-of-band light, so that we can reduce our exposure times and the resolution of the camera, and operate the IRIS host vehicle at highway speeds.

- Fabrication of the on-chip masking scheme of Figure 5. One of the main sources of noise and sensitivity to vibrations in the current IRIS prototype is the way that the masking scheme has been aligned with the CCD chip. We can eliminate these problems by fabricating the alternate-row mask of Figure 5 and mounting it directly on the CCD chip. Our efforts to convince chip manufacturers such as Texas Instruments to collaborate with us in this direction have not been fruitful thus far. Therefore, we now trying to do this ourselves in the Nanofabrication facility at UCLA, where all the necessary equipment is already in place. Using specialized handling equipment under a microscope, we will remove the cover of the CCD, attach and align the mask, and replace the cover. To avoid problems of misalignment, we will use a special adhesive that hardens only when exposed to ultraviolet light. This way, we will perform the alignment under red lights, and when everything has been set to meet our specifications, we will flood the device with ultraviolet light. This on-chip masking scheme will make our prototype much smaller than it currently is, and it will allow us to significantly reduce the sensor's sensitivity to vibrations, as well as the required exposure times.

- Implementation of the IRIS-2 configuration in hardware, and modification of the ranging algorithms to utilize the geometry of Figure 3. This will endow the IRIS sensor with stereoscopic vision and it will eliminate the need for any a priori knowledge about the target, such as the distance between the taillights.

- Integration of the IRIS sensor with a low-cost radar sensor. The combination of two different sensor technologies with complementary capabilities will allow us to produce a new generation of highly reliable sensor systems capable of operating under nearly all weather conditions and without any infrastructure requirements.

We are also in the process of filing patents for this technology, and we continue to pursue possible collaborations with industrial partners for its commercialization.

6. DISSEMINATION AND PUBLICITY

In order to disseminate our results and attract attention to this project, we have taken the following steps:

1. We prepared and presented a poster at the TRB booth of the NAHSC Exposition Hall during Demo 97 held in San Diego, California, on August 7-10.

2. We have developed a Website (http://ansl.ee.ucla.edu/iris) for the project that contains on-line information and publications.

3. We presented seminars on this technology at

   - the ITS Enabling Technologies Symposium organized by the John L. Volpe National Transportation Systems Center of the U.S. Department of Transportation in Cambridge, Massachusetts, on September 16, 1997,

   - the Center for Autonomous Systems of the Royal Institute of Technology in Stockholm, Sweden, on January 30, 1998,

   - industry meetings held at UCLA in February and April 1998, which has resulted in initial contacts with Fuji Corporation, and

   - Rockwell International in El Segundo, California, in August 1998.

4. We have written a paper entitled “Intelligent Sensors and Control for Commercial Vehicle Automation,” which was presented at the IFAC Workshop on Advances in Automotive Control, held in Mohican State Park, Ohio, February 27–March 1, 1998; the paper also appeared in the *Proceedings* of the Workshop.

5. Our project has recently received attention from the media; it was featured in the cover story (The Cutting Edge) of the Business section of the *Los Angeles Times* on January 12, 1998; the cover page of that article is shown in Figure 18.

MONDAY TECHNOLOGY SPECIAL

SECTION D

MONDAY

JANUARY 12, 1998 CC1

BUSINESS

THE CUTTING EDGE

Los Angeles Times

Road Scholars

IRIS SCHNEIDER / Los Angeles Times

From left, researcher Phyllis Nelson, student Buddy Nelsen and UCLA professors Oscar Stafsudd and Ioannis Kanellakopoulos.

'Smart' cars that could revolutionize transportation and reduce pollution are moving closer to reality with technology being developed in the Southland.

By KAREN KAPLAN, TIMES STAFF WRITER

For most people, the thought of doubling the number of cars on the Southland's crowded highways inspires dread and maybe a little road rage. But for a group of research engineers along the Tech Coast, it is something of a Holy Grail.

These researchers at UCLA, USC, UC Irvine and UC Riverside are developing intelligent transportation technologies that, when fully implemented, could boost freeway capacity as much as 200% while reducing travel time and cutting back on air pollution.

"Certainly there's been a whole sequence of technological innovations since Henry Ford, but the thing that hasn't changed since Henry Ford's time is the interaction between the vehicle and the roadway," said Steven Shladover, deputy director of Partners for Advanced Transit and Highways, a research effort administered by UC Berkeley that receives about $6 million a year from the California Department of Transportation, as well as funding from other sources.

PATH's goal is to make cars "smarter" by enabling

Please see CARS, D6

Traffic Control

"Smart" cars could increase the traffic capacity of today's freeways without requiring any new road construction. Sensors and actuators would work in tandem to pack more cars in traffic lanes. The devices would help vehicles maintain a safe distance from one another and keep them cruising at a constant speed, eliminating much of the abrupt acceleration and braking that causes traffic jams and adds to air pollution.

How It Works

Researchers are designing intelligent cruise-control systems that allow one car to follow another. The high-tech car will automatically slow down if the car ahead slows, or speed up if the lead vehicle accelerates. Tiny sensors attached to the front of a car can measure the distance between it and the vehicle ahead. Engineers at UCLA are using infrared technology to build an inexpensive system that uses lasers like those found in compact disc players.

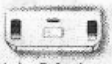

Capturing Reflections

A laser bounces off reflective patches in the taillights and reflective paint in the license plate. These reflections become data that are recorded and "compared."

Electronic Snapshots

A microprocessor instructs a charge-coupled device, or CCD, to take an electronic picture (brown arrow) of the car in front. A split second later, a laser fires and the CCD camera takes another picture (pink arrow), this time with the laser illuminating reflective patches on the taillights and the license plate.

Gauging the Distance

The CCD repeats this process over and over, comparing the position of the license plate and taillights, which enables it to determine the relative speed of and distance between the two cars.

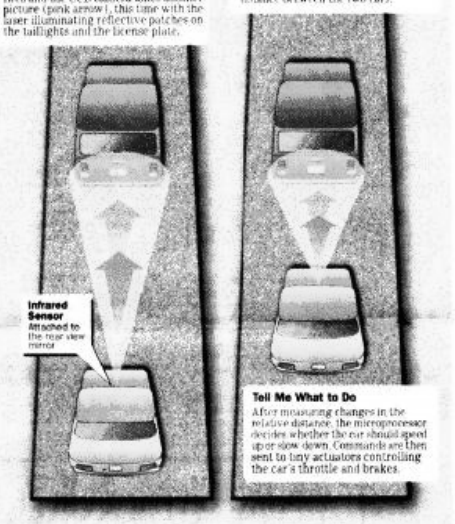

Better Technology, More Cars

As more drivers adopt "smart" technology, highway capacity grows. Vehicles per hour, per lane

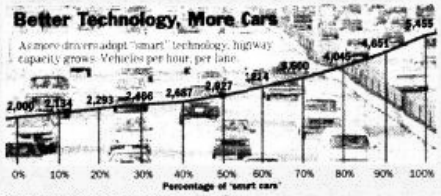

Source: Center for Advanced Transportation at USC, UCLA

MATT MOODY / Los Angeles Times

Figure 18

The front page of the Los Angeles Times article featuring our IRIS sensor.