ITS INSTITUTE
Intelligent Transportation Systems

Integration of Infrared Imaging for a Head Up Display Lane Keeping and Collision Avoidance System

Final Report

Prepared by:

Pi-Ming Cheng
Craig Shankwitz

Intelligent Vehicles Lab
Department of Mechanical Engineering

Intelligent Transportation Systems Institute
Center for Transportation Studies

CTS 08-09

CENTER FOR TRANSPORTATION STUDIES
UNIVERSITY OF MINNESOTA

Technical Report Documentation Page

1. Report No. CTS 08-09	2.	3. Recipients Accession No.	
4. Title and Subtitle Integration of Infrared Imaging for a Head Up Display Lane Keeping and Collision Avoidance System		5. Report Date June 2008	
		6.	
7. Author(s) Pi-Ming Cheng, Craig Shankwitz		8. Performing Organization Report No.	
9. Performing Organization Name and Address Intelligent Vehicles Lab Department of Mechanical Engineering University of Minnesota 1100 Mechanical Engineering 111 Church Street SE Minneapolis, Minnesota 55455		10. Project/Task/Work Unit No. CTS project # 2004056 11. Contract (C) or Grant (G) No.	
12. Sponsoring Organization Name and Address Intelligent Transportation Systems Institute University of Minnesota 200 Transportation and Safety Building 511 Washington Avenue SE Minneapolis, Minnesota 55455		13. Type of Report and Period Covered Final Report 14. Sponsoring Agency Code	
15. Supplementary Notes http://www.cts.umn.edu/pdf/CTS-08-09.pdf			

16. Abstract (Limit: 200 words)

A driver assistive system, which uses high accuracy, differentially corrected GPS (DGPS), high-accuracy geospatial databases, radar, computers, and driver interfaces (both a Head Up Display (HUD) and a tactile seat), has been developed to help a driver maintain lane position and avoid collisions during periods of low visibility. These systems have been tested and deployed both in Minnesota and in Alaska.

Collision avoidance information is provided to a driver through the HUD. Objects located within the HUD field of view determined to be a threat to a driver are indicated as square boxes. White boxes represent an advisory, and if the detected object is fewer than fifty feet or three seconds from a collision, the box turns to red (a warning).

Drivers who have used the system have speculated that a more accurate representation of the object projected on the HUD would make the driving task easier. Emerging Super and High Dynamic Range Cameras (SDRC and HDRC, respectively) appeared to be a feasible, inexpensive means with which to address drivers' concerns.

This report documents the research program. SDRC and HRDC technology failed to meet expectations, but infrared imagery was successfully integrated with the standard HUD.

17. Document Analysis/Descriptors Visibility Enhancement System, Head Up Display, Infrared sensors, Thermal imaging		18. Availability Statement No restrictions. Document available from: National Technical Information Services, Springfield, Virginia 22161	
19. Security Class (this report) Unclassified	20. Security Class (this page) Unclassified	21. No. of Pages 28	22. Price

Integration of Infrared Imaging for a Head Up Display Lane Keeping and Collision Avoidance System

Final Report

Prepared by:

Pi-Ming Cheng
Craig Shankwitz

Intelligent Vehicles Lab
Department of Mechanical Engineering

Intelligent Transportation Systems Institute
Center for Transportation Studies

June 2008

Published by:

Intelligent Transportation Systems Institute
Center for Transportation Studies
University of Minnesota
200 Transportation and Safety Building
511 Washington Ave SE
Minneapolis, MN 55455

Acknowledgements

The authors would like to thank Mn/DOT for its continued support of IV Lab Projects, and for the care and fuel it provides the IV Lab for the use of the Safeplow. The Safeplow is an invaluable research tool that allows a number of new technologies to be evaluated in a safe fashion, both on the test track and on roadways. Without this support, technologies and applications like those described herein would never reach a workable level of development.

This support is greatly appreciated.

This project was sponsored by the US DOT RSPA through the ITS Institute at the University of Minnesota.

Table of Contents

List of Tables

List of Figures

Executive Summary

The Minnesota Department of Transportation (Mn/DOT) is responsible for ensuring that state roadways remain clear and safe during the harsh winter months. Fortunately, a recent emphasis on the application of scientific and engineering principles in the winter maintenance industry has led to the development of more effective equipment to aid in this task.

In 2001, operational testing of a driver assistive device for low visibility snowplowing conditions began in Minnesota, and continues today in both Minnesota and Alaska on a small fleet of vehicles. The primary focus of this testing remains snow removal equipment, but the system has been installed and used on state patrol cars and Airport Rescue and Fire Fighting (ARFF) equipment as well.

This driver assistive system uses high accuracy, differentially corrected GPS (DGPS), high-accuracy geospatial databases, radar, computers, and driver interfaces (both a Head Up Display (HUD) and a tactile seat) to help a driver maintain lane position and avoid collisions during periods of low visibility. As a driver travels along a road for which a geospatial database (aka, digital map) exists, the driver is provided a conformal, augmented "virtual" view of the landscape in front of the vehicle, even under conditions where visibility out of the windshield may be limited to a few feet. These low visibility conditions occur frequently during the snow removal process.

The fidelity with which the virtual view through the Head Up Display depends on the amount of detail placed in the geospatial database and the amount of data provided by on-vehicle sensors. Prior to this research project, collision avoidance information was provided by on-board radar. This radar provides range, range rate (i.e., relative velocity), and azimuth angle to the identified targets. The radar indicates "something" is "out there," and where "it" is, but radar cannot identify what "it" is. As such, radar information is provided to the driver through an icon projected by the HUD. The objects are presented as squares, rendered in white outlines if the object is more than fifty feet from the vehicle, or more than three seconds from a collision. This provides an advisory to a driver that an object is ahead. The radar target icon is projected in the proper perspective; as the target approaches the host vehicle, the icon increases in size as it is projected on the HUD.

Should the object being tracked by the radar be determined to be closer than fifty feet or less than three seconds from a collision, the icon presented to the driver moves from a white rectangle to a red one, indicating a warning condition.

As a result of operational testing, drivers requested more information from the radar targets. Drivers typically determine following distance based on the type of vehicle ahead. Smaller vehicles typically can brake at higher rates than heavy trucks; therefore, snowplow operators tend to provide more headway to smaller vehicles. However, this size information is unavailable from the radar sensor.

To provide a driver with additional information, imaging sensors were explored. At approximately the same time the IV Lab driver assistive system was first-tested [1.], General

Motors introduced their Night Vision option for Cadillac automobiles [2.]. This night vision system, developed with Raytheon, used a passive IR sensor to identify thermal gradients (i.e., “warm” animals on the “cold” road) in the landscape to identify possible threats to drivers. Images were provided to a driver via a HUD, but the images were not conformal; i.e., the projected images are not aligned with the view through the windshield. Moreover, this technology was proprietary to GM, and not available for general consumption. The technology has since been discontinued by GM because of small sales volumes [3.].

Concurrent with the emergence of the IV Lab HUD and the Cadillac Night Vision was the announcement of a new imaging technology: Super Dynamic Range Cameras, or SDRC. The promise of SDRC is that it offers both greater sensitivity in the near IR range and higher dynamic range than conventional CCD imaging devices. Greater sensitivity in the near IR range provides the capability to detect thermal gradients; high dynamic range allows the camera to not “wash out” images in the presence of ambient light sources (i.e., approaching headlights). Moreover, the sensors were expected to be relatively inexpensive, on the order of $2,000 in small quantities.

However, the promise of SDRC was not achieved. No vendor in the US could provide a unit for either sale or demonstration; one vendor from Italy claimed to have SDRC capability, but that vendor was unable to demonstrate adequately that the camera offered near IR sensitivity needed for this application.

Even without an SDRC image source, a second research topic remains: can video information provided by an imaging source be integrated into the synthesized virtual display provided by the HUD? The HUD information is synthesized from vehicle position, heading, database information, and radar sensor information, and is projected in the proper perspective enabled by the optical properties of the HUD system. The image data provided by the thermal imager has to be transformed (in real time) both in terms of perspective and resolution to the HUD coordinate reference frame. A method to perform these transformations has been developed and is executed using a conventional thermal imaging device. This process and its results are described herein.

Project findings are that SDRC cameras have yet to be produced which meet the goals of high dynamic range, near IR sensitivity, and affordability for vehicle applications. However, the technique of integrating image data with synthesized images based on geospatial databases, radar, and vehicle position and heading information has been developed and demonstrated. Should SDRC devices which meet original specifications become available, they can be readily integrated into the existing IV Lab driver assistive device.

Chapter 1
Introduction

In 2001, operational testing of a driver assistive device for low visibility snowplowing conditions began in Minnesota, and continues today in both Minnesota and Alaska on a small fleet of vehicles. The primary focus of this testing remains snow removal equipment, but the system has been installed and used on state patrol cars and Airport Rescue and Fire Fighting (ARFF) equipment as well.

This driver assistive system uses high accuracy, differentially corrected GPS (DGPS), high-accuracy geospatial databases, radar, computers, and driver interfaces (both a Head Up Display (HUD) and a tactile seat) to help a driver maintain lane position and avoid collisions during periods of low visibility. As a driver travels along a road for which a geospatial database (aka, digital map) exists, the driver is provided a conformal, augmented "virtual" view of the landscape in front of the vehicle, even under conditions where visibility out of the windshield may be limited to a few feet. These low visibility conditions occur frequently during the snow removal process.

The fidelity with which the virtual view through the HUD depends on the amount of detail placed in the geospatial database and the amount of data provided by on-vehicle sensors. Prior to this research project, collision avoidance information was provided by on-board radar. This radar provides range, range rate (i.e., relative velocity), and azimuth angle to the identified targets. The radar indicates "something" is "out there," and where "it" is, but radar cannot identify what "it" is. As such, radar information is provided to the driver through an icon projected by the HUD. The objects are presented as squares, rendered in white outlines if the object is more than fifty feet from the vehicle, or more than three seconds from a collision. This provides an advisory to a driver that an object is ahead. The radar target icon is projected in the proper perspective; as the target approaches the host vehicle, the icon increases in size as it is projected on the HUD.

Should the object being tracked by the radar be determined to be closer than fifty feet or less than three seconds from a collision, the icon presented to the driver changes from a white rectangle to a red one, indicating a warning condition.

As a result of operational testing, drivers requested more information from the radar targets. Drivers typically determine following distance based on the type of vehicle ahead. Smaller vehicles typically can brake at higher rates than heavy trucks; therefore, snowplow operators tend to provide more headway to smaller vehicles. However, this size information is unavailable from the radar sensor.

To provide a driver with higher fidelity information, image information could be integrated with the information provided by both the radar sensors and the on-board positioning and database systems. The radar provides information where obstacles are in terms of range and azimuth angle, but doesn't provide any information as to "what" the obstacle is. The goal of the research is to integrate image information with sensor and on-board information to provide a driver information regarding both the "where" and the "what" associated with nearby vehicles.

The integration of imaging, sensor, and on-board information was motivated by the introduction of High Dynamic Range Cameras (HDRC) and Super Dynamic Range Cameras (SDRC). The promise of these technologies was greater sensitivity in the near infrared region and higher dynamic range, producing less sensitivity to glare and oncoming headlights. Moreover, these new imaging devices were to cost much less than conventional thermal imaging sensors (normally known as infrared cameras). HDRC cameras were expected to cost approximately $2,000 in small quantities; SDRC cameras, offering lower dynamic ranges, cost approximately $500.

A three step process lead to the integration of the imaging sensors with the visibility enhancement system. Step one consisted of identifying candidate sensors, and procuring those which were demonstrated to offer the performance necessary to provide a driver improved fidelity of the image projected by the HUD. The second task was to characterize the sensor in terms of actual near-IR sensitivity, resolution, and dynamic range. The third task was to develop a means to "blend" the image from the imaging sensor with the virtual image generated by the integration of position, radar, and geospatial data available to the IV Lab driver assistive system. This blending requires image compression, eyepoint coordinate transformation, and field of view alignment so that images properly overlay.

It should be noted that the promise of SDRC and HDRC cameras were not realized in this study. Neither SDRC or HDRC cameras offer sufficient near IR sensitivity nor sufficient dynamic range to detect small temperature gradients or respond to oncoming headlight glare. Despite these shortcomings, a technique to integrate image information with the synthesized information projected in the IV Lab HUD was developed, and tested using true thermal imaging equipment. As will be shown, even expensive (~USD$15,000) thermal imaging cameras exhibited difficulty in the driving environment. Although near IR sensitivity was sufficient for the application, thermal imagers also suffered from headlight glare because of limited dynamic range. Regardless, a technique to integrate IR with the standard HUD system was developed and demonstrated; should improved SDRC or HDRC cameras become available, image information can be readily integrated into the HUD.

In the sequel, the three aforementioned steps in the integration process are described, with results documented with both still images and video captured during testing on Minnesota Trunk Highway 7 in Excelsior and Minnetonka, MN.

Chapter 2
Sensor Procurement and Characterization

Three classes of imaging sensors were considered: SDRC cameras, HDRC cameras, and true thermal imagers.

2.1 Evaluation of SDRC Cameras

SDRC cameras are most prevalent. Panasonic, considered a world leader in surveillance cameras, produces a line of SDRC cameras, known as Super Dynamic Range 2 Cameras. SDRC 2 cameras are based on a ¼" CCD (Charge Coupled Device), and provide a dynamic range of 36 dB. The IV lab procured a model (WV-CL920V), and performed an evaluation. The cost of the SDRC 2 camera, with lens, was approximately $500. Typical specifications for a Panasonic SDRC 2 CCD camera are shown below in Table 1.

Table 1. Performance specifications for Panasonic SDRC II CCD Imaging Sensors. Note the dynamic range of 36 dB.

Effective Pixels	768 (H) x 494 (V)
Scanning Area	1/4" type Super Dynamic II CCD
Horizontal Scanning Frequency	15.734 kHz
Vertical Scanning Frequency	59.94 Hz
Signal-to-noise Ratio	50 dB (AGC off, weight on)
Dynamic Range	36 dB (selectable on/off)
Minimum Illumination	2 lx (0.2 fc)
Video Output	1.0 v[p-p] NTSC composite / 75 Ohm

Evaluation of the Panasonic WV-CL920V camera was performed at the intersection of US Highway-52 and County Highway 9 in Goodhue County, MN. The Panasonic WV-CL920V camera is equipped with electronic enhancement and claims sensitivity up to 32 times greater than that of a standard CCD camera. When combined with a near infrared light source, the camera can also be used for near infrared surveillance in settings with no visible light. However, the problem arising with vehicles at night is that vehicle headlights spread a beam of light on the road; this complicates the automatic detection process. For the SDRC camera and illuminator to work for this purpose, the supplemental illuminator needs to produce an ambient light brighter than the headlights of vehicles. When the camera was tested along with an infrared illuminator at night, it was found that the illumination provided was insufficient to robustly detect vehicles in the image. Figure 1 and Figure 2 show images captured with the Panasonic WV-CL920V camera at night. It is noted that the only blobs of vehicle headlights and headlight reflection on the road can be seen. With this quality of image, image processing software cannot reliably determine whether the blob constitutes one or two vehicles (Figure 1) or whether two blobs arise from one or two automobiles, one automobile and two motorcycles, or four motorcycles (Figure 2). The inability to resolve the number and types of vehicles in low ambient light conditions renders the SDRC technology inadequate for HUD augmentation.

Figure 1. A captured image from a Panasonic WV-CL920V camera. It is difficult to determine if there were one or two vehicles in a big light blob.

Figure 2. A captured image from a Panasonic WV-CL920V camera. It is difficult to determine if two light blobs belong to one or two vehicles or motorcycles.

2.2 Procurement of HDRC Cameras

Identifying a source for HDRC Cameras proved rather difficult. The early promise of a relatively inexpensive sensor, which prompted this investigation, never was manifest in the procurement of an affordable, available, demonstrable device.

After a significant search, the only manufacturer who produces a camera which could be considered in the HDRC class is Neuricam, from Trento, Italy. The specifications for the Neuricam NC-5100 HDRC camera are provided in Table 2 below.

Table 2. Performance specifications for Neuricam NC-5100 CMOS Imaging Sensors. Note the claimed dynamic range of 120 dB.

Effective Pixels	640 (H) x 480 (V)
Scanning Area	Grey Scale CMOS Sensor
Horizontal Scanning Frequency	16.67 KHz
Vertical Scanning Frequency	26.04 Hz
Signal-to-noise Ratio	Not Specified
Dynamic Range	120 dB
Minimum Illumination	Not Specified
(Video) Output	Ethernet, Serial RS-232, and 3 optoisolated I/O lines

The Neuricam HDRC camera offers a claimed 120 dB dynamic range, significantly more than the Panasonic claim of 36 dB.

Procurement of this camera was difficult, not only because of the language issues between Minnesota and Italy, but also because Neuricam was unwilling to provide a unit for demonstration.

> Dear Mr. Cheng,
>
> Unfortunately we do not give cameras for evaluation purposes. An evaluation purpose is actually the most common purpose of our customers, who buy the cameras in order to investigate how to use our components in a more complex systems they develop. For doing this they invest around 2.000 Euro.
>
> What I can do, is to allow you a 5% discounts, which is usually given to distributors only.
>
> If you decide to buy the camera NC5100-SDK for 1.890 Euro + around 50 Euro shipping costs, please give me a confirmation, we will then communicate our bank coordinates for the payment.

Thank you

Gualtiero Chini

Pi-Ming Cheng wrote:

Hi,

After a discussion with our chief engineer, we are wondering if it's possible that we can get a demo unit of a NC5100 to try out for a week or so to determine if it meets our needs. We'll be willing to pay for all (international) shipping/insurance costs.

Please understand that 200 Euros (or USD$2650) for a camera is quite a lot. We'll be more comfortable to spend the money if we can try it out first.

Thank you very much and look forward to hearing from you!

=> The NC5100 is the version of EtherCam equipped with the CMOS
=> monochromatic logarithmic sensor (named NC1802 Pupilla), that we
=> produce. Being logarithmic it assures a high dynamic range (120 dB),
=> therefore it is the most proper for avoiding saturation and to perform
=> better with changing lighting conditions.
=>
=> Please find attached the brochure of the EtherCam and a picture.
=> In the following mail you will find the brochure of the Pupilla sensor.
=>
Gualtiero Chini
Marketing & Sales

NeuriCam S.p.A.
Via Santa Maria Maddalena 12, 38100 Trento, Italy
Tel +39-0461-267 777; fax +39-0461-260 617
e-mail: chini@neuricam.com; http://www.neuricam.com

Because Neuricam was unable/unwilling to provide a demonstrator before a purchase order was issued, an on-line demonstration was requested. Neuricam reluctantly agreed, and made available a web site where the response of the camera to a variety of lighting conditions could be observed. Unfortunately, this demonstration proved to be of little value. First, the actual lighting conditions were difficult to either quantify or qualify because the images provided could not be compared to the actual conditions as sensed by the human visual system. The human eye is a high performance imaging system, and can serve as a reference for other imaging systems. However, to serve as a reference, it has to be exposed to the same conditions as the sensor being evaluated. This is impossible over the web.

Second, the images provided by Neuricam were not as impressive as what were expected. The ability of the camera to adjust from very bright to dark conditions (i.e., a condition which occurs commonly when a vehicle approaches and passes in the on-coming lane) was slow. Second, Neuricam was hesitant to provide video data showing the response of the NC-5100 to oncoming and passing headlights. Even with the claimed 120 dB dynamic range, based on the video information provided by Neuricam, images are likely to be washed out when headlights are directed at the camera.

In summary, the quality of the images from the Neuricam NC-5100 were of insufficient to justify an investment of $2650 for a camera clearly not capable of providing near IR sensitivity and sufficient dynamic range to avoid image "washout" in the presence of oncoming headlights.

2.3 Evaluation of IR Cameras

Because the project has a second objective, the integration of actual image information with the virtual image synthesized by the IV Lab driver assistive system, a camera capable of providing sufficient sensitivity in the near IR range and sufficient dynamic range (or conversely, low sensitivity in the visible light range) is needed. After reviewing product specifications and convincing suppliers to provide demo units, three IR cameras (FLIR A20, Indigo Omega, and Electrophysics PV320-A2) were selected for evaluation. In this evaluation, the objective was to select a camera with the highest performance vs. cost ratio (with performance valued more than cost).

Figure 3 shows images from the three infrared cameras that were evaluated. All three cameras were focused on roughly the same area. The white pixels in the images indicate objects with high temperature and the black pixels indicate objects with low temperature. All three cameras have the same image resolution (160x120 pixels), but from the comparison of the images, it is clear that A20 offers the highest image clarity.

The IR camera selected for the development project was a FLIR A20. This is a true thermal imager, capable of detecting thermal gradients in highway environments. It is important to note that the FLIR A20 has a resolution of 160 x 120 pixels, or 1/16 the resolution of a regular NTSC camera (640x480 pixels). Specifications for the FLIR A20 Thermal Imaging Camera are provided below in Table 3.

Blending a 160x120 pixel thermal image with the synthesized 640x480 image produced by the IV Lab Driver Assistive System provides a significant technical challenge, especially when the horizontal and vertical fields of view of the thermal imager (25° x 19°, respectively) are different than those of the Head Up Display (35° x 18°, respectively) onto which the images are blended. However, if the "blending" can be done with disparate image resolutions and fields of view, implementation with a SDRC or HDRC (should units of sufficient quality appear) will be straightforward, by comparison.

Figure 3. Images from the three IR cameras evaluated.

Table 3. Performance specifications for FLIR A20 Thermal Imaging Sensor.

Effective Pixels	160 (H)x 120 (V)
Scanning Area	Focal plane array (FPA) uncooled microbolometer
Spectral range	7.5 to 13 µm
Temperature ranges	-20° C to +250° C (-4 to +482° F)
Thermal sensitivity @ 50/60Hz	0.12° C at 30° C
Accuracy (% of reading)	± 2°C or ± 2%
Video output:	RS170 EIA/NTSC or CCIR/PAL composite video
FireWire/Ethernet output	8/16-bit monochrome, 8-bit color
Available Lenses: (FOVs,/Minimum Focus Distance)	• 12° Telescope: (12° x 9°/1.2m), • 25° Normal: (25° x 19° / 0.3 m), • 45°Wide angle: (45° x 34°/0.1m)

Chapter 3
Sensor Integration

3.1 Driver Assistive Display

A Driver Assistive Display is a high fidelity graphical system capable of providing an accurate visual representation of the roadway and surroundings local to the vehicle. The system integrates data from DGPS, a high accuracy geospatial database, and radar sensors. The result is a "see-through" highly accurate Head Up Display (HUD) based image, displaying road boundaries with the proper perspective, relevant geographic features (e.g., signs, bridges, jersey barriers, etc.), and obstacles impeding the motion of the vehicle.

However, information from radar suffers two fundamental shortcomings. The first is its inability to detect and "follow" obstacles moving perpendicular to the vehicle direction of travel. In contrast, the human visual processor is very sensitive to cross track motion. The human operator, when presented with visual information in the HUD provided by an IR source, will detect cross track motion of pedestrians and animals which would have likely been missed with radar. The second shortcoming is that automotive radar will provide information pertaining to *where* an obstacle is (range, range rate, and azimuth angles), but not information regarding *what* the obstacle is. Infrared sensors can provide information regarding *what* is detected, but unless stereo sensors are used, cannot provide accurate information regarding *where* an obstacle is. Clearly, imaging and radar technologies complement one another.

3.2 IR Image Integration

In this study, the information from the IR camera was fused with all the information synthesized by the driver assistive system and displayed on the HUD. A FLIR A20 IR camera was purchased and installed on the Safeplow. A Pelco EH3512-1 camera enclosure equipped with a heater and a blower was used to house the IR camera to protect the camera from wide ranges of temperature, humidity, precipitation, and sunlight. Because regular glass blocks the passage of infrared light, the camera enclosure was fitted with a special IR transparent crystal that is suitable for outdoor applications. Figure 4 shows the Safeplow with the FLIR A20 camera on the roof. The IR camera was installed directly on top of the driver's seat so that the field of view of the camera can be physically aligned with driver's viewpoint, eliminating one coordinate translation calculation.

Figure 4. FLIR A20 IR camera on the Safeplow

Images from the FLIR A20 IR camera are captured using a PIXCI SV5 imaging board (http://www.epixinc.com/products/pixci_sv.htm). The specifications of the video capture board are provided in Table 4 below.

Table 4. Specifications for PIXCI SV5 imaging board.

Connections	- 4 Pin DIN: S-Video Input - Two BNC-Jacks: Composite Video Inputs - DB15: TTL I/O Triggers - Cables optional
Video Input	- Color or Monochrome Video Acquisition: S-Video, RS-170, CCIR, NTSC, PAL - Resolution-Pixels: 754x480: RS-170, NTSC, S-Video 922x580: CCIR, PAL, S-Video - Resolution-Depth: 8 bit: RS-170, CCIR YUV [4:2:2]: NTSC, PAL YCrCb: S-Video - Capture/Display Rate: 30 fps: RS-170, NTSC, S-Video 25 fps: CCIR, PAL, S-Video
Data formats to PCI Bus	- Monochrome 8 bit - YCrCb (UYUV, YUV4:2:2): 16 bit
Capture Rate	- 30 fps: RS-170, NTSC, S-VIDEO - 25 fps: CCIR, PAL, S-VIDEO
Transfer Rates	- Requires a PCI motherboard with burst mode to host memory data rates of at least 30 MB/S.
Display - Windows	- Display resolution as per installed VGA device driver. - A DCI compatible S/VGA adapter is required for real-time display.
Bus Requirements	- 32 bit PCI Bus Master slot - 0.55 Amps @ +5 Volts - 4.913 in. by 3.350 in.

The PIXCI XV5 board utilizes 128MB of frame buffer memory allocated during the computer boot up process. The frame buffer memory is divided into several different image buffers so that the computer software can analyze one captured video frame while the next video frame is being captured into a different buffer in order to achieve real-time video capturing.

Once an image is in the frame buffer, an image processing subroutine known as the Canny algorithm [4.] is carried out to extract objects of interests (mainly vehicles) from the thermal image. The Canny algorithm is a multiple-step algorithm designed to detect a wide range of edges in an image. An image is processed in the following order:

(1) **Reducing noise:** A raw image is convolved with a Gaussian mask and the output is a lightly blurred version of the original. This step is to ensure that a single noisy pixel has little effect on the Gaussian smoothed image. It softens the image, but reduces the presence of any spurious pixels, improving the probability that the line detection process will be robust.

(2) **Finding the intensity gradient of the image:** An edge in an image may point in any direction. The Canny algorithm uses 4 convolution masks to compute the intensity gradient in the horizontal, vertical, and diagonal (up left and up right) directions. The convolution result for each mask direction for each pixel is stored. After the gradients for each of the four directions is computed, each pixel is evaluated to determine the value and direction of the gradient with the greatest magnitude. This step creates a map of the maximum intensity gradient and its direction for each pixel in the original image

(3) **Tracing edges through the image**: The higher intensity gradients are more likely to be edges. However, it is possible that a given gradient switches from not being on an edge to being an edge. The Canny algorithm uses thresholding with hysteresis in four steps to determine the continuity of a potential edge.

   a. **Step 1**: **Initial thresholding**. A gradient above a certain level is likely an edge. The first step is to identify those pixels associated with the high gradient as likely elements of an edge.

   b. **Step 2: Non-maxima pixel suppression.** The pixels identified in step 1 are checked to see whether each pixel above the threshold has a value greater than its two neighbors along the direction associated with that threshold. If it is the local maximum, its value kept; otherwise, the gradient is assigned a value of zero.

   c. **Step 3: Bi-thresholding.** With the result of Step 2 available, two more thresholds (t1 and t2, where t1<t2) are applied, producing images T1 and T2, respectively. T2 will have less noise, fewer false edges, but more edge gaps than T1. The Canny algorithm is most sensitive to these thresholds. For a class of images, i.e., highway thermal imaging, a iterative process is needed to determine proper values for t1 and t2.

   d. **Step 4: Edge linking.** In image T2, edge segments are linked to produce continuous edges. At the end of an edge in T2, image T1 is searched for edge segments which bridge the edge gaps in T2. This produces the desired binary, edge-detected image.

The algorithm is able to follow a faint section of a given line but avoids identifying a few noisy pixels that do not comprise a line. Once the process is complete, the result is a binary image where each pixel is marked as either an edge pixel or a non-edge pixel.

It is to be noted that the Canny algorithm used in the study was adapted from the Intel Open Source Computer Vision Library (OpenCV Library) [5.] . The OpenCV Library has many subroutines that can be applied in real-time computer vision applications. The OpenCV Library is an open-source library and can be freely incorporated into other software.

After the objects of interest are extracted from a thermal image, they are correlated with radar return information and projected with the proper perspective on the HUD. Because the infrared images have a lower resolution (160x120 pixels) than the HUD (640x480 pixels), only a portion of the HUD is overlaid with the IR image. This both preserves the objects of interest in the correct dimensions and locations and prevents image distortion during the transformation and projection process. Figure 5 shows the flowchart of the IR image integration onto the HUD.

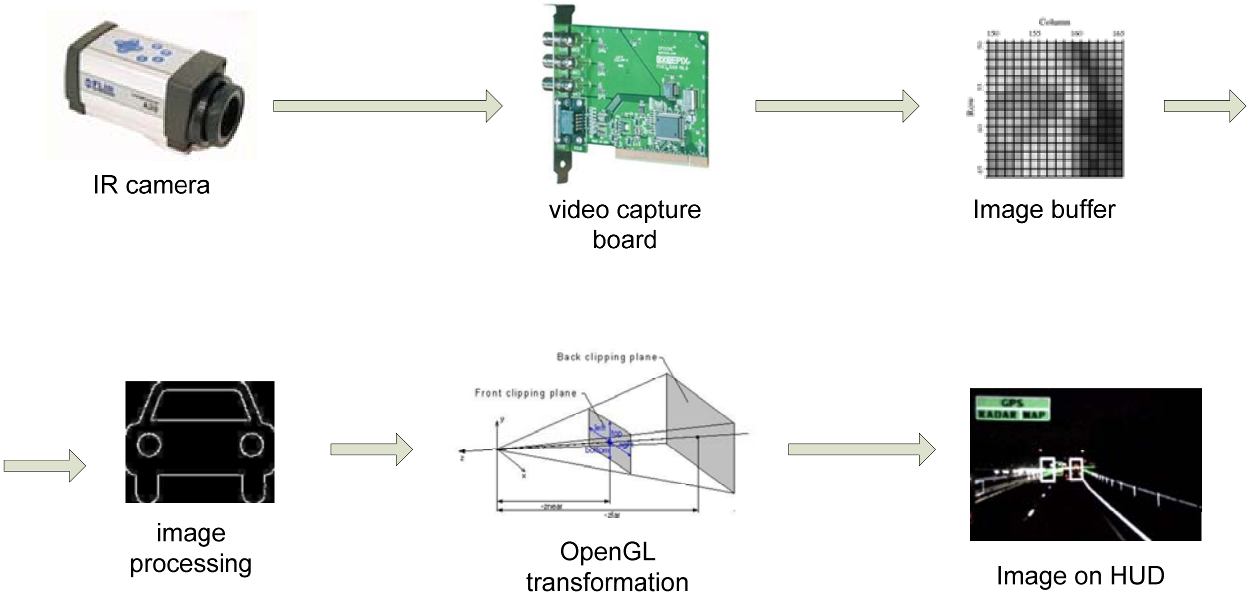

Figure 5. Flowchart of the IR Image integration onto the HUD.

3.3 Results

This section provides representative still images captured during testing on Minnesota Trunk Highway 7 between Excelsior and Minnetonka, MN. Video (and a Windows Medial Player Codec) are provided on the CD attached to the back cover of this report.

Figure 6 shows the system tracking a vehicle with an IR image superimposed onto the HUD. The white blob is from the Canny edge detection algorithm and indicates an object of interest on the IR thermal image. The rectangle box is from radar return information.

Figure 6. System tracking a vehicle with an IR image superimposed onto the HUD. The white blob indicates an object of interest.

Figure 7 shows the system tracking multiple vehicles with an IR image superimposed onto the HUD. Due to vehicle lights and distance to tracked vehicles, white blobs on vehicles that are farther away from the Safeplow are not very clear.

Figure 7. System tracking multiple vehicles with an IR image superimposed onto the HUD.

Background noise on the IR image in Figure 7 exists due to vehicle lights and the relatively close environmental temperature between the road and the objects of interest (tracked vehicles). This causes the image processing subroutine to not function as well as expected, resulting in the overlap between the outputs of the Canny algorithm and the background pixels. A background noise filter was added to the image processing routine to further enhance the captured thermal images. The background noise filter removes all image pixels that are below or near the ambient temperature in a captured thermal image.

Figure 8 shows the system tracking two vehicles with the additional background noise filter. Please note that the background noise is almost nonexistent in Figure 8.

Figure 8. System tracking two vehicles with the background noise filter activated.

Figure 9 shows another picture of the system tracking multiple vehicles with the additional background noise filter.

Figure 9. System tracking multiple vehicles with the background noise filter activated.

Chapter 4
Conclusions and Recommendations

This study successfully integrated thermal imagery into the driver assistive system and provided a stand-alone obstacle display interface. The driver assistive system can now synthesize the HUD information from vehicle position, heading, database information, radar sensor information and thermal imagery and project images onto the HUD in the proper perspective. However, the FLIR A20 IR camera has a price tag of ~USD$15,000. At such a high price, it's unlikely that the technology can be widely deployed. Moreover, the FLIR A20 IR camera has an output image resolution of 160x120 pixels, which is only 1/16 of the resolution of the HUD (which has a resolution of 640x480 pixels). In order to preserve the proper perspective on the much higher resolution HUD, the processed IR imagery can only cover part of the HUD screen. The thermal image also suffered from quality degradation in the graphics transformation stage and "washed out" in the presence of headlight glare at night because of limited dynamic range.

The technologies that enable HDRC and SDRC are still under development. These technologies, if successful, will offer both greater sensitivity in the near IR range, and higher dynamic range than conventional CCD images. Greater sensitivity in the near IR range provides the capability to detect thermal gradients; high dynamic range allows the camera to not "wash out" images in the presence of ambient light sources (i.e., approaching headlights). More importantly, these sensors are expected to be sold (per unit) at less than 1/10 of the price of an IR camera. In the near future, should the HDRC or SDRC which meet their original specification become available, the software and techniques developed and demonstrated in this study can be readily integrated into the existing driver assistive system.

References

[1.] Heon-Min Lim, Bryan Newstrom, and Max Donath. "A Conformal Augmented Head Up Display for Driving under Low Visibility Conditions," Proceedings of the 5th International Symposium on Advanced Vehicle Control, Ann Arbor, MI, August 2000.

[2.] http://www.gm.com/company/gmability/safety/avoid_crash/new_features/night_vision.html

[3.] http://www.ivsource.net/modules.php?name=IV_Archives&file=article&sid=93

[4.] J. Canny, *A Computational Approach To Edge Detection*, IEEE Trans. Pattern Analysis and Machine Intelligence, 8:679-714, 1986.

[5.] http://www.intel.com/technology/computing/opencv