A Contribution to Laser Range Imaging Technology

NASA Contract Final Report

RICIS Preface

A Contribution to
Laser Range Imaging Technology
Nasa Contract Final Report

Principal Investigator:
Professor Rui J. P. deFigueiredo
Dept. of Electrical and Computer Engineering
Rice University, Houston, Texas 77251-1892

Prepared with help from graduate research assistant
Bradley S. Denney

Date: February 4,1991

1 Overview

The goal of the project was to develop a methodology for fusion of a *Laser Range Imaging Device* and camera data. Our initial work in the project led to the conclusion that none of the LRID's that were available were sufficiently adequate for this purpose. Thus we spent the time and effort on the development of a new LRID with several novel features which elicit the desired fusion objectives. In what follows, we describe the device developed and built under the contract. Funds from other sources were also used in the implementation of this task (such as grants from the Texas Advanced Technology Program and Texas Instruments).

The Laser Range Imaging Device (LRID) is an instrument built by us under the contract which scans a scene using a laser and returns range and reflection intensity data. Such a system would be extremely useful in scene analysis in industry and space applications. The LRID will be eventually implemented on board a mobile robot.

The current system has several advantages over some commercially available systems. One improvement is the use of X-Y galvonometer scanning
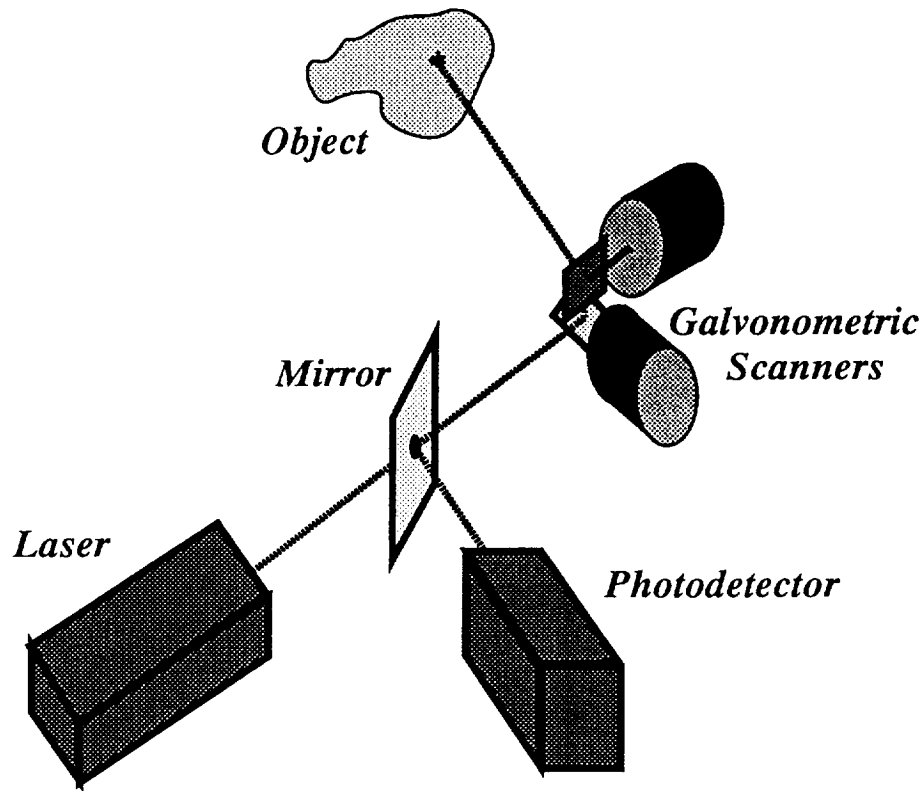
Figure 1: The LRID system developed

mirrors instead of polygonal mirrors present in some systems. The advantage of the X-Y scanning mirrors is that the mirror system can be programmed to provide adjustable scanning regions. For each mirror there are two controls accessible by the computer. The first is the mirror position and the second is a zoom factor which modifies the amplitude of the position parameter. Another advantage of the LRID is the use of a visible low power laser. Some of the commercial systems use a higher intensity invisible laser which causes safety concerns. By using a low power visible laser, not only can one see the beam and avoid direct eye contact, but also the lower intensity reduces the risk of damage to the eye, and no protective eyeware is required. (Note: In applications where eye safety is not a concern it is logical to use a higher powered infra-red laser). In addition the use of the visible laser facilitates the alignment of the system's optics.
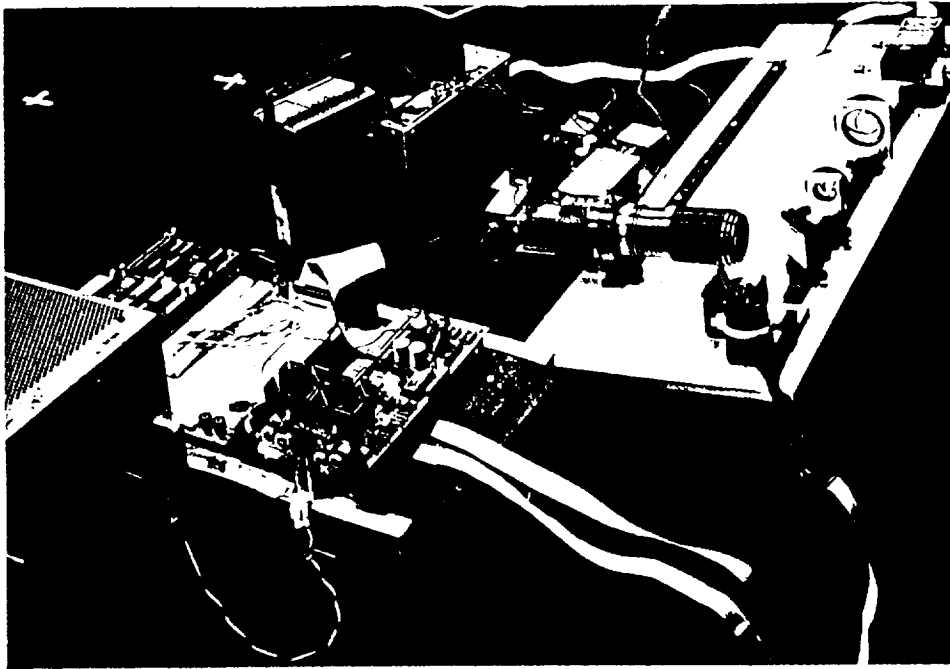
Photo 1:   The LRID experimental system. Here we see the Single Board Computer and our scanner controller (bottom left) which drives the X-Y scanners on our optical array (top right). Along side the optical setup is the circuitry for the transmitter, receiver, and phase detection (top left).

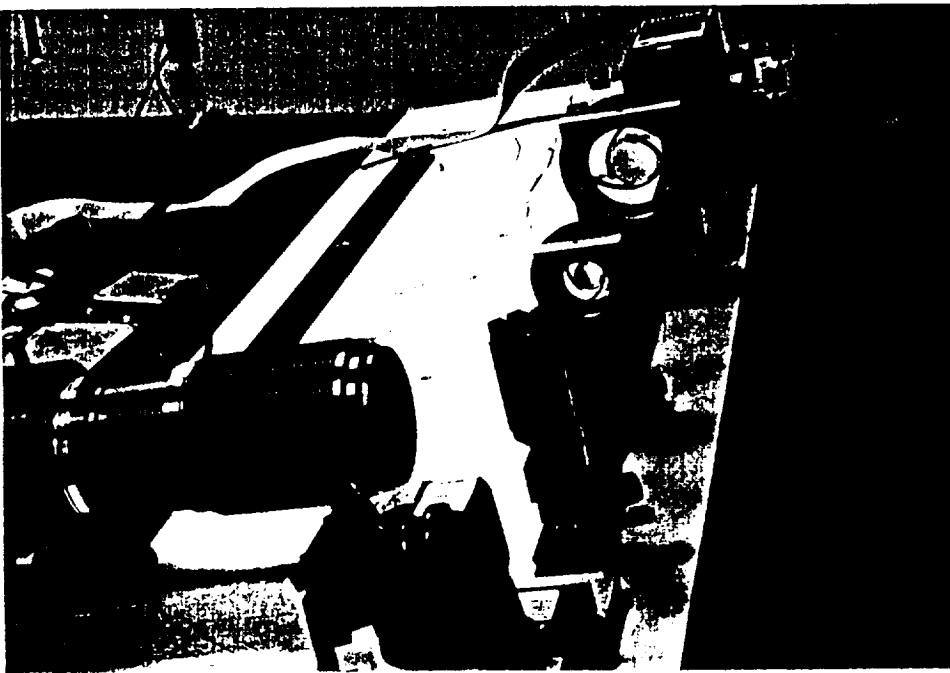Photo 2:   The LRID optical Array. Starting from the top right corner, a red (670nm) laser beam is emitted from a diode laser. The Beam passes through two lens and through a small hole in a mirror and then is directed on the scene by two galvonometer scanning mirrors (bottom right). The reflected beam is returned back to the photo-detector through a standard camera lens (left) via the three mirrors.

The LRID (see Figure 1 and Photos 1 and 2) obtains range data by modulating the output beam intensity of the laser. The laser being used in our system is a red (670nm) diode laser with a maximum output of 5mW. The diode laser is used due to its ability to be easily modulated. Simply modulating the current to the diode causes intensity modulation. The modulated beam is then focused and passed through a small hole in the back of a mirror and onto the X-Y scanning mirrors. A portion of the laser light is reflected back to the scanner mirrors and then to the front side of the mirror with the small hole. Since the optical aperture of the mirrors is about 5mm and the hole is 1mm, the loss due to the hole is much less than that which would have occurred using a beam splitter. The return light from this mirror is then focused onto an avalanche photo diode sensor.

The object distance is determined by comparing the phase of the detected reflection intensity with the original modulation signal. The phase shift represents the distance traveled by the beam. In the LRID system a modulation frequency of 30.4 MHz is used. Since our beam travels at the speed of light, the wavelength of the intensity modulation is given by $\lambda = c/f$. Thus, the maximum distance traveled corresponding to a phase shift of 360° is λ. Note that there is some ambiguity in the range detector due to the fact that a distance traveled of d will look the same as a traveled distance of $d + n\lambda$ Where n is any non-negative integer.

The distance from the scanner to the object is one half of the round trip distance. If we assume that we will only be able to detect objects within a distance of $\lambda/2$ from the laser, our phase shift, θ, will be restricted to be from 0 to 360°. Thus using the phase shift data and given the modulation frequency, f, one can determine the distance to an object by

$$d = \frac{\theta}{360°}\frac{c}{2f}$$

where c is the speed of light.

Currently the LRID is using a two quadrant phase detector which will be enhanced to cover all four quadrants in the future. With a four quadrant detector our ambiguity interval will be 4.93m. One possible enhancement to this system would be be the use of a second modulation frequency. In this case the exact same optical setup could be used. Only the transmitter and receiver need be partially modified. By mixing two different modulation frequencies, one could first detect a rough range reading from the target, and

then use this knowledge to determine which ambiguity interval the object is actually in. Once this interval is known a more accurate range measurement can be calculated.

Besides the collection of range data, the LRID also collects intensity data. This data is time averaged intensity data from a certain point in the object. The intensity data corresponds to the diffuse reflectance of the object point in question at the laser wavelength and the angle of incidence of the beam with the object (the relationship can be given by Lambert's cosine law). A red filter at the sensor will block out most of the room noise. While this intensity image is similar to camera data, it has some distinct advantages. For example since the image is obtained from active illumination (i.e. the laser beam) there is no need to worry about poor lighting. And since the return reflection to the sensor and the illumination beam are coaxial there will be no shadows in the image. This type of image is ideal for further processing. The drawback of using this system for intensity images is that the the frame acquisition rate is rather slow.

The scanning system is in fact the major time consumer of the system. While the scene scan time of about 800ms (for 128 by 128 pixels at a scan of $\pm 20°$ in both the x and y direction) is comparable to that of the Odetics system, it is much slower than video rates. This limits the practicality of using the LRID as a *stand alone* real time imaging system. However, by using data from other sources such as a camera, the LRID can be programmed to look at some subset of the scene which may require fewer scan lines or a smaller mirror swing. This would reduce the scan time as well as unwanted data.

2 Description of Scanning Modes

The LRID can be operated in various scanning modes which facilitate its use in a more global setting. Standard raster scans are available to collect samples over the full range of the scanning device. However other (less time consuming) raster sensors (such as video cameras) may provide information which reduce the number of data points needed. Thus we have spent a great deal of time on the treatment of points of interest.

Given a point list we have developed a fixed time *vector scan*. This scan moves from point to point according to a previously unknown point list (i.e.

the point list is continuously being updated with more points). Since the points are few with respect to a raster scan, the desired result may be more quickly obtained.

3 Utilizing LRID Data

Using range data one can estimate various parameters of the object being looked at. Some of these basic parameters include range and direction from scanning unit (position), orientation of a surface, return intensity, and diffuse reflectance at the laser wavelength.

The first quantity, position, is directly measurable from the phase detector and the direction of look of the scanner (e.g. 10° up, 8.3° left, and 1.23 meters away, with respect to the scanner). A quick judgment of range of an object of interest can be quickly made by looking at a point on the object with the scanner and taking a range measurement.

One should note that the range measurement is ambiguous and could be actually

$$r_{\text{actual}} = r_{\text{measured}} + n\lambda_m,$$

where n is some positive integer. In other words there's no way of telling if the phase shift is φ or $\varphi + n360°$. This problem can be taken care of by using multiple modulation frequencies. The longer wavelengths would determine which ambiguity interval the object is in and the shorter wavelength would make more precise measurements of the range. This, of course, requires additional hardware and a minimal amount of processing.

The next measurement, orientation at a point, is estimated by looking at the points neighbors to obtaining the gradient of the point. Again quick measurements can be made by looking at a small number of points (three points are sufficient) in the neighborhood of the point of interset.

The returned intensity is a useful quantity. As explained earlier, The LRID coaxial illumination/reflection path make the laser scanned images free from shadows and poor lighting flaws.

The characteristic diffuse reflectance of an object can be estimated by realizing that the averaged received radiant flux is proportional to the averaged transmitted radiant flux with a proportionality constant of,

$$\bar{F}_R = \frac{\alpha A_R \rho_d \cos\theta}{\pi r^2} \bar{F}_T$$

where α is the transmission of the optics, A_R is the scanner receiving area, ρ_d is the diffuse reflectance of the object at λ, θ is the angle of incidence of the beam on the object surface, and r is the range from the scanner to the object point. Using our measurement of r and $\bar{F}_R$, and our estimate for the orientation (which is related to θ), we can find an estimate for ρ_d. This quantity will vary from material to material and could be used to add another dimension to object classifier algorithms.

4 Results

The results of the system built can be seen in the photographs 3-5 attached to this report. Due to the use of a two quadrant phase detector we were limited to an ambiguity interval of $\lambda_m/4$. Where λ_m is the modulation wavelength of the system. The noise associated with the results is mainly due to two factors.

The first source comes from using a small signal strength. since the return signal is proportional to $1/r^2$, there are only a few photons that are actually sensed by the receiver. In addition black targets have a very low photon return especially at large angles of incidence. While a low signal strength (low power laser) is used in the laboratory for safety reasons, it can be increased for field applications for use in applications (especially when eye safety is not a concern). It has been suggested that in some applications, the objects of interest can possess passive optical elements to help with their detection. For example an object could be equipped with strategically located retro-reflectors which give a very high directed return.

Since we also have an intensity image generated form the LRID, we can tell which pixels suffered from low return levels. We can then ignore that data.

Another source of noise comes from the receiver design. While the design is generally effective, it is not able to handle the wide range of incoming signal levels. We have recently proposed a different design which should deal effectively with this problem. In addition the new design contains a four quadrant detector which will allow us to detect objects in the full range.

5   Fusion with Camera

The idea of sensor fusion comes from the fusion of data generated from multiple sensors to a single data set. In our case, we have two very different sensors which provide different types of data. There are different levels at which fusion can occur. The lowest level of fusion deals with the physical devices, the middle level deals with the fusion of the raw image data, and the highest level is concerned with the fusion of data at the scene understanding level.

5.1   Low Level Fusion

At the lowest level the fusion is physical to the devices, in other words the two devices are combined physically as one device to give us just one sensor. This level however, is not readily realizable.

It is also possible to create new types of sensors combining the two sensors. One example is a profile generator. Using the laser (ignoring the ranging capabilities), one can scan an image using scan lines and make some profile measurements with the camera when the camera and the LRID are not close together.

5.2   Mid Level Fusion

The lowest practical level (the middle level) which utilizes the full capabilities of the sensors is the fusion of the raw data sets of the two sensors. In our case the data is transformed into a similar axis representation and combined into some usable form for example, the data could be put in a list of 3-D points. Each point in the list, would represent a point on the surface of some object in the scene, and would also contain associated parameters of that point, such as reflected passive illumination, and reflected active illumination. In another scheme a 2-D point system could be used, where each point is associated with its range from the LRID, and both its reflected passive and active illumination.

5.3 High Level Fusion

The fusion can of course come at an even higher level and that is at the level of image understanding. Due to the time considerations (i.e. the LRID is slow compare to the camera when obtaining entire images), we can view sensor fusion as more of a decision method concerning which sensor to use.

Using a camera image one can come up with regions of interest. These regions are related to the application. The desired solution to the problem at hand may not be available using a single camera or may require an enormous amount of computation.

In an example such as navigation, there may be some previous knowledge of the scene. The camera can be used (due to its speed), to verify the existence of certain objects in the scene and give a general idea of the object's direction from the sensors. At this point, the LRID can be used to look at *interesting* subsets of the scene and return range data of the objects as well as reflectance properties which may be used to verify the model.

5.4 Coordinate Transformation

It is worth noting that coordinate transformations can be easily verified. When using a laser beam and camera which have overlapping light frequencies, the coordinate system transformations can be verified by looking at the spot created by the laser with the camera at fixed beam locations.

6 Conclusion

The development of the LRID has provided us with a useful tool with some unique features that enhance its use in a sensor fusion environment. The LRID's scanning modes allow rapid vector type scanning which is valuable when concerned with regions of interest much smaller than the entire raster. In addition the LRID's variable raster scan window size allows different degrees of resolution in the images obtained. Also with the use of multiple modulation frequencies, one can make A wide range of range measurements with fairly good accuracy.

Using the LRID which we developed under the contract, we now have a useful tool for implementing some of the ideas outlined in the previous section. Probably the best fusion of the LRID and camera data would come
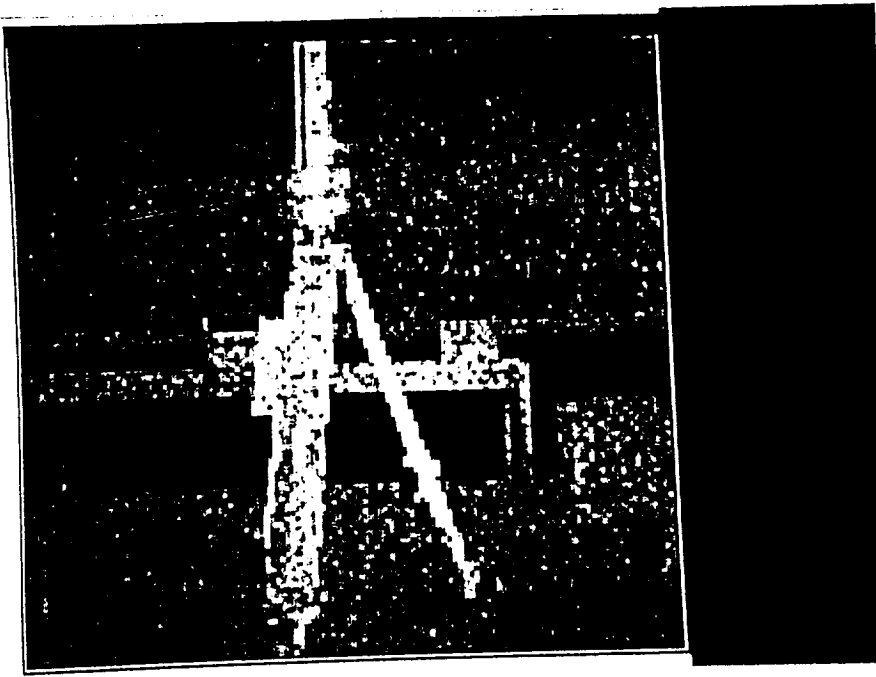
Photo 3: Preliminary System Results. Picture here is a 128 by 128 pixel image of the scene shown in Photo 6. The effects of the current two quadrant phase detector can be seen. Darker grey levels represent pixels which are closer to the laser scanner. The white pixels are near the end of the interval and the pixels start to become darker again as one continues out from the scanner.
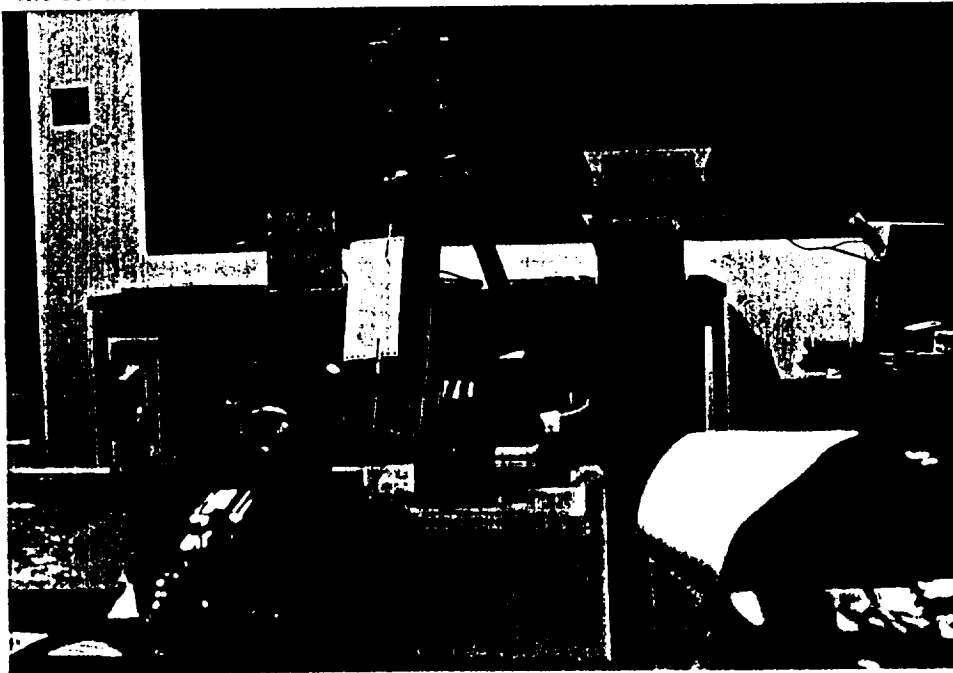

Photo 4: Actual scene in the Rice Robotics Lab imaged above by the LRID. (Note that camera and scanner positions are not identical).
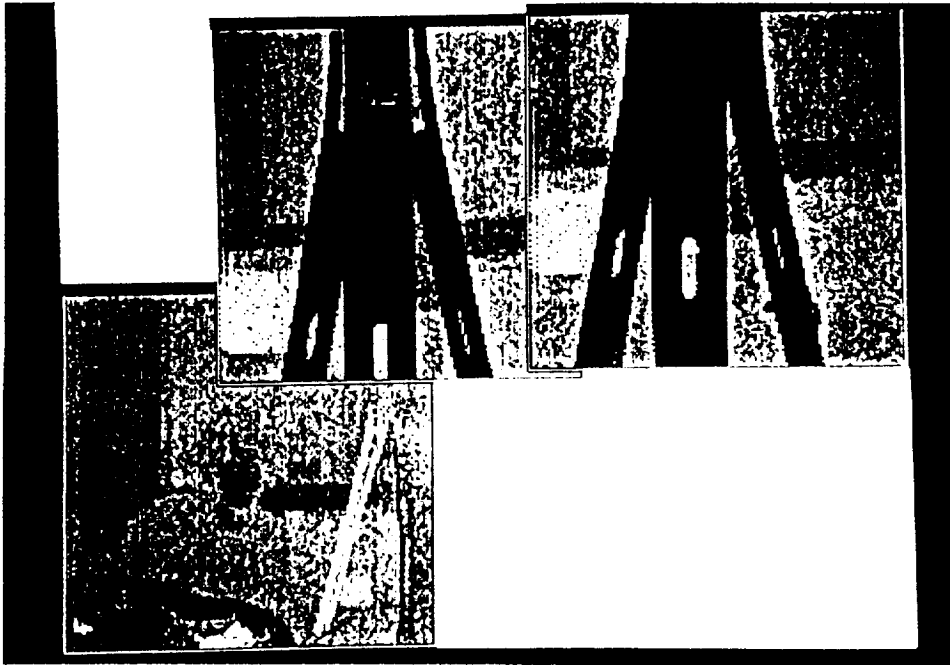
Photo 5:   Several Images aquired from the LRID.