U.S. Department
of Transportation

**National Highway
Traffic Safety
Administration**

People Saving People
http://www.nhtsa.dot.gov

DOT HS 808 506

October 1996

Interim Report

Development of Performance Specifications for Collision Avoidance Systems for Lane Change, Merging, and Backing

TASK 5: Crash Countermeasure Technology Investigation

FHWA

Technical Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
DOT HS 808 506		

4. Title and Subtitle	5. Report Date
Development of Performance Specifications for Collision Avoidance Systems for Lane Change, Merging, and Backing. Task 5 Interim Report: Crash Countermeasure Technology Investigation	October, 1996
	6. Performing Organization Code
	8. Performing Organization Report No.
7. Author's) Philip Moffa, Joseph Austin, G. Sam Dow, Ike Ikizyan and Mark Hibben	
9, Performing Organization Name ond Address TRW Space and Electronics Group One Space Park Redondo Beach, CA 90278 USA	10. Work Unit No. (TRAIS)
	11. Contract or Grant No.
	13. Type Of Report and Period Covered Task Order Study October 1996 through July 1996
12. Sponsoring Agency Name and Address National Highway Traffic Safety Admin. (NHTSA) Office of Collision Avoidance Research (OCAR) 400 Seventh St. SW, Room 5301 Washington, D-C. 20590 USA	14. Sponsoring Agency Code NRD-51

15. Supplementary Notes

*F04606-90-D-0001/0052 (TRW-Air Force) under Air Force-NHTSA Inter-agency Agreement DTNH22-93-X-07022 (Air Force-NHTSA)

16. Abstract

An investigation of the technologies that are important for crash countermeasures against lane Change and merging accidents has been conducted. The sensor technologies have been limited to radar and lidar because of the relatively long ranges required for high closing velocity accident avoidance. The state of the art of digital signal processors has also been determined. The other key technology, namely, display technology is being investigated in other studies.

The preliminary performance specifications that were developed earlier as part of this program have been utilized as a benchmark which the capability of each technology has been measured against. For sensor technologies, the ability to detect the specified targets over the defined volume in space is a key performance goal. Equally important is the immunity from both false and nuisance alarms. Issues concerning public safety from radiation exposure, and mitigation of interference have also been addressed.

In the area of processing, a set of requirements for processor speed, architecture, and memory was derived from an estimate of the computations needed to accomplish the detection of targets and the further calculations necessary to determine their speeds and discriminate against non-threatening or inappropriate targets.

The results indicate that both radar and lidar technology are capable of performing the detection function and both have enough waveform flexibility to mitigate against interference. Both can be utilized in a manner consistent with the prevailing safety standards and both have solid state implementations that could lead to low-cost, highly reliable components when the demand becomes more real. Processors are also advancing a rapid rate. Small low cost processors capable of performing the necessary computations in a timely manner already exist and progress continues towards faster and cheaper units. The overall cost of the crash countermeasure system cannot yet be determined; for one thing, it will vary in accordance with its level of crash avoidance capability.

17. Key Words	18. Distribution Statement
Collision Avoidance Systems, Lane Change, Merging, Technology Investigation, Radar, Lidar, Laser, Processor	This document is available to the public through the National Technical Information Service, Springfield, VA 22161

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of Pages	22. Price
Unclassified	Unclassified		

Form DOT F 1700.7 (8-72) Reproduction of completed page authorized

TABLE OF CONTENTS

ACRONYMS

A/D	analog to digital
AM	amplitude modulation
ANSI	American National Standards Institute
APD	avalanche photodiode
c	speed of light
ca-cfar	cell-averaged constant false alarm rate
CAS	collision avoidance system
cfar	constant false alarm rate
dB	decibels
dBsm	decibels relative to one square meter
DSP	digital signal processor
DVI	driver vehicle interface
FCC	Federal Communications Commission
FFT	Fast Fourier Transform
FIFO	first-in first-out
FM	frequency modulation
FMCW	frequency modulated continuous wave
FOV	field of view
FWI-IM	full-width at half-maximum
8	acceleration due to gravity
GHz	gigahertz (one billion per second)
go-cfar	greater-of constant false alarm rate
HBT	heterojunction bipolar transistor
HEMT	high electron mobility transistor
IC	integrated circuit
IEEE	Institute of Electrical and Electronics Engineers
IFOV	instantaneous field of view
ITR	infinite impulse response
kHz	kilohertz (one thousand per second)
LNA	low noise amplifier
MBE	molecular beam epitaxy
MAC	multiply-accumulate
MFLOPS	millions of floating point operations per second
MHz	megahertz (one million per second)

MIPS	millions of instructions per second
MMIC	millimeter wave monolithic integrated circuit
MMW	millimeter wave
MOPS	millions of operations per second
MPE	maximum permissible exposure
MPET	maximum permissible exposure time
mph	miles per hour
NHTSA	National Highway Traffic Safety Administration
os-cfar	ordered-statistic constant false alarm rate
Pd	probability of detection
P fa	probability of false alarm
P fd	probability of false detection
PM	pseudomorphic
PN	pseudo-random noise
prf	pulse repetition frequency
RBW	receiver bandwidth
RAM	random access memory
RF	radio frequency
rms	root mean square
ROM	read-only memory
R&D	research and development
SAE	Society of Automotive Engineers
scr	signal to clutter ratio
snr	signal to noise ratio
SV	subject vehicle
TBD	to be determined
TBR	to be reviewed
v c o	voltage controlled oscillator

SUMMARY

An investigation of the technologies that are important for crash countermeasures against lane change and merging accidents has been conducted. The sensor technologies have been limited to radar and lidar because of the relatively long ranges required for high closing velocity accident avoidance. The state of the art of digital signal processors has also been determined. The other key technology, namely, display technology is being investigated in other studies.

The preliminary performance specifications that were developed earlier as part of this program have been utilized as a benchmark which the capability of each technology has been measured against. For sensor technologies, the ability to detect the specified targets over the defined volume in space is a key performance goal. Equally important is the immunity from both false and nuisance alarms. Issues concerning public safety from radiation exposure, and mitigation of interference have also been addressed.

In the area of processing, a set of requirements for processor speed, architecture, and memory was derived from an estimate of the computations needed to accomplish the detection of targets and the further calculations necessary to determine their speeds and discriminate against non-threatening or inappropriate targets.

The results indicate that both radar and lidar technology are capable of performing the detection function and both have enough waveform flexibility to mitigate against interference. Both can be utilized in a manner consistent with the prevailing safety standards and both have solid state implementations that could lead to low-cost, highly reliable components when the demand becomes more real. Processors are also advancing a rapid rate. Small low cost processors capable of performing the necessary computations in a timely manner already exist and progress continues towards faster and cheaper units. The overall cost of the crash countermeasure system cannot yet be determined; for one thing, it will vary in accordance with its level of crash avoidance capability.

Crash Countermeasure Technology Investigation
(Task 4.5 of CET)

1.0 Scope

The state-of-the-art of the various technologies relevant to lane change/merge collision avoidance systems (CAS) has been investigated. See reference (1) and references therein for a more general description of crash avoidance technologies. For each technology, both available systems and those requiring further development were evaluated. For purposes of this study, we limited our technology time horizon to the year 2000 so that systems utilizing them could be realized by 2005 on the appropriate platforms.

The technologies were measured against the yardstick represented by the preliminary performance specifications set forth in Task 4.4 (2) of Phase I for the functional goals related to a lane change/merge system as specified under Task 4.2 (3). Issues such as false alarm rates, detection probabilities, cost, reliability, safety, platform compatibility, etc. were addressed, treating the CAS as a system with interfaces to the platform and to the humans who must accept its usage.

The sensing component and the processing component of the CAS are often intertwined. Hardware design was traded against processing complexity. The measures of effectiveness being cost, risk and performance. Although performance was the driving requirement, cost and technology risks were also included in this evaluation. This sensor-processor interaction was addressed mostly at the top level, such as the processing complexity associated with each sensor type. In some instances, the interaction delved deeper into sensing schemes within each hardware type, such as in the case of FMCW versus pulse-Doppler radars.

Since only radar and lidar-based CAS can reliably meet the long range requirements for high closing speed lane change/merge CAS, they were the only ones addressed in this sensor technology evaluation.

Factors concerning the safety of the public with regards to the radiated energy, and the potential for interference mitigation when these systems are deployed in large numbers were considered. This effected the operational frequency of the system and the total amount of power that can be safely radiated. For radar systems, the choice of frequency is constrained by its interaction with the atmosphere and the practical requirement to keep the antenna(s) as small as possible. There are formal safety standards on the total power that can be radiated which are essentially independent of frequency. For laser systems, frequency selection is again constrained by atmospheric propagation issues, by background radiation reduction, and even more importantly by eye safety considerations.

For this study, our emphasis was on the requirements for CAS in the near-term. Accordingly, we did not investigate vehicle control technology, off-board remote sensing and roadway communications. Moreover, display issues are being addressed by separate NHTSA-sponsored work at VRTC and Battelle (4). We will defer to results from that study,

In Section 2, our approach to this technology investigation will be discussed. Both the sensing and processing issues will be addressed. This section concludes with a reproduction of the preliminary performance specifications for a lane change/merge CAS that were presented in the Task 4.4 Interim report (2).

Section 3 contains the technology assessment for the sensing portion of the CAS and Section 4 presents the processing technology investigation. Section 3 begins with general remote sensing considerations and also discusses the key sensor functions of detection, range measurement, and speed and bearing determination. The appendix will elaborate on the theory of detection. Finally, conclusions are presented in Section 5.

2.0 Approach

Our Phase I study of lane-change, merge and backing collisions culminated in the specification of preliminary performance requirements for collision avoidance systems (CAS). These preliminary performance specifications would allow the CAS to attain functional goals which represent changes to the crash situations that would have prevented-the accidents. While we have performed a technology survey in Task 4.3 (5) by testing existing off-the-shelf hardware, the emphasis has been a top-level assessment of current state-of-the-art. We had made no attempt at delving into sensor phenomenology beyond characterizing the detection of the systems empirically with standard size targets in static and dynamic situations. We had no knowledge of the target detection and discrimination algorithms a priori. Whatever insights we gained were through empirical observation and deduction. As for the preliminary performance specifications, we have arrived at the results through simulations and analysis and have largely undertaken the effort without worrying about how well the existing systems will meet them. In the current task, we started with the preliminary performance specifications of Task 4.4 as the requirements and addressed the question as to whether they can be met with advanced technologies, within the time horizon that we have chosen, We built on the preliminary survey that we had started in Task 4.3.

2.1 Sensing Technologies

As stated under 1 .O, only radar and lidar technologies will be considered. The driving requirement is detection range, with the smallest target being a pedacyclist for a lane change/merge CAS. Our sensor technology investigation was centered around detection performance, given a required detection zone (range and field of regard). The detection range is in general a function of power, aperture, receiver noise temperature, target characteristics and system losses. System losses include attenuation due to weather, hardware losses, system inefficiencies and processing losses. These losses are in general sensor specific and some are design dependent. Some system parameters are limited not simply by technology but also by safety considerations (e.g. power level), by platform constraints (e.g. aperture size) and by field of regard considerations (e.g. aperture placement). Target characteristics are sensor-dependent and target detectability depends not just on the targets themselves but also on the ability of the sensor to distinguish them from the background clutter. Our study made an assessment as to what is technologically feasible given these phenomenology and system constraints.

In Task 4.3, we noted that most existing systems are short-range (less than 20 ft) and that most can detect the smallest and most stressing targets (child and pedacyclist), at least under benign (clutter-free), static conditions. Our emphasis in this task was to examine system behavior at longer ranges, as are found necessary from our Task 4.4 investigations to mitigate against fast-closing collisions.

A key requirement for CAS is to minimize false and nuisance alarms. False alarms are spurious detections caused by noise in the system, nuisance alarms are true detections of objects that are not considered a threat. The noise figures and signal to noise ratios characteristic of each technology determine the false alarm rate. Nuisance alarms are caused by a variety of conditions, such as, large objects in the sidelobes of the radiation patterns, and small non-threatening objects in the main radiation beam. To avoid nuisance alarms, the reaction of the various sensors to a variety of objects typically encountered but usually non-threatening (especially during turns) must be evaluated. These include parked cars, trees, guardrails, speed bumps, some curbs, opposing traffic, etc.

In addition, the sidelobe levels attainable from radar antennas, and optical lenses were determined. Detection of objects in radiation sidelobes is also a function of the dynamic range of the targets encountered. If a large object like a truck has significantly more scattering return than say a small child, then it can be readily detected even when only a very small fraction of the radiated power illuminates it. The dynamic range of the scattering returns from all viable targets and their variations with orientation, weather conditions, and coverings needed to be ascertained in each relevant wavelength band.

While sensor phenomenology helps discriminate clutter, it is often insufficient as man-made clutter can often mask the target signature. Target discrimination algorithms are devised to eliminate unwanted clutter using a priori knowledge of the target and clutter themselves. These algorithms are an integral part of any sensor system design and often cannot be divorced from it. However, we made an artificial demarcation here and examined them thoroughly in Section 4, Processing Technologies.

2.2 Processing Technologies

Under this area, we divided our study into three steps. First, we examined the processing requirements generic to each type of sensor up through the point that a target signature can be obtained. Based on these considerations, we drew broad conclusions about processing complexity. Second, we investigated target discrimination algorithms, alluded to in Section 2.1 above. Under this category are all processing required prior to the issuance of a warning. The driving requirement here is that false positives must be minimized while true positives must be maximized. Finally, we compared these requirements against state-of-the-art processors and determined if they are technologically feasible, &en cost, size and power constraints.

Because of the short-range nature of the existing systems, the target discrimination algorithms are not overly stressed, geometry having limited the clutter scene itself to one that is manageable. As the target range and/or field of regard of the system is increased, the algorithms will be stressed. We started with the performance of existing systems and extrapolated their performance as necessary to the required range and field of regard found in Task 4.4. Again the most stressing targets have to be considered against the largest size clutter objects. Such a situation may be encountered, for example, in the case of a motorcycle target at maximum range versus a large truck in the side-lobe of a radar system (but outside the intended field of regard of the CAS) at short range.

The test of existing systems in Phase I indicated that most existing systems are short range as the collisions resulting from lane-change, merge and backing have been largely perceived to be a proximity type collision, at least to first order. In the few cases of systems with ranges in excess of 20 feet, discrimination algorithms which claim to work against ground clutter are not very robust. These systems are mainly radar systems where velocity measurements can be made relatively easily, using for example, a pulse-Doppler radar. Ground clutter are eliminated using the Doppler shift from the target return as an added discriminator. These ground clutter have Doppler shifts centered around the vehicle velocity. For a 60 mph vehicle, this shift is on the order of 20 (15) kHz for a 94 (77) GHz radar operating frequency. The efficacy of these types of algorithms were examined. The hardware complexity of a pulse-Doppler radar versus an FMCW system was also explored. Their inherent capability in velocity measurements were examined.

Often, the discrimination algorithms employed have hardware implications. While velocity measurements using radars are best performed using a pulse-Doppler radar design, a simpler radar design can be used to measure velocity by successive range measurements, although at degraded accuracy, increased processing and increased system latency. These disadvantages must be traded against hardware complexity/cost. These factors were examined in our study.

Conversely, high performance and sophisticated algorithms may require very complex hardware. This is the case for coherent imaging radars whereby 2-D radar images can be formed for target identification and discrimination. These techniques are widely employed in military applications but are unlikely to be realizable in the near-term for commercial use.

We see the clutter discrimination algorithms as a critical technology for implementing longer range systems than those that exist today. While those for some lane-change CAS do exist, they have only limited performance.

2.3 Driver Vehicle Interface Functional Requirements and Inputs

The CAS can be conveniently divided into three main subsystems. The sensor subsystem provides the ability to detect the potentially dangerous situation, the processor analyzes the data generated by the sensor, and the driver vehicle interface (DVI) converts the processor outputs into information easily accessible to the driver. The state of technology for each of these subsystems is being investigated. The processor and sensor technologies have been presented here and the DVI technology is being researched in a parallel effort. (4)

In addition to understanding the overall performance specifications for the CAS, these requirements need to be flowed down to the subsystem level. In addition, the interfaces between the subsystems must be defined. Since the sensor and processor are intimately linked and since they are being considered here and the DVI elsewhere, the interface between the combined sensor and processor, and the DVI will be presented now.

The interfaces between the vehicle, the combined sensor and processor, and the DVI are represented in Figure 2.3.1. A general approach has been taken in drawing this figure in that the performance specifications for the lane change/merge CAS have not been finalized and hence the interfaces are also uncertain. For example, the interface between the vehicle and the CAS sensor/processor includes turn signal status, steering inputs, vehicle speed, and gear specification. The turn signal status will be utilized either for CAS activation or for modifying the display modality. The two most likely implementations of a lane change/merge CAS (and the ones most often utilized in the systems we tested) are the one where the CAS is activated by the turn signal (or some other indication of a lane change/merge maneuver), and the one which is always on when the vehicle is in forward gear. For the always-on type, the warning level for the presence of an object in the designated zones around the vehicle would be non-intrusive such as a visual one. It would only become more intense (auditory or haptic) when the turn signal was on or there was some other indication (steering, lane tracking) that a lane change/merge maneuver was occurring. Steering inputs may be utilized to predict that a lane change or merge may be occurring or that an inadvertent drifting is taking place. The speed is a basic input for the CAS. For example, it may be used so that the processor can discriminate against fixed objects which would be measured to be either approaching or receding at the speed of the instrumented

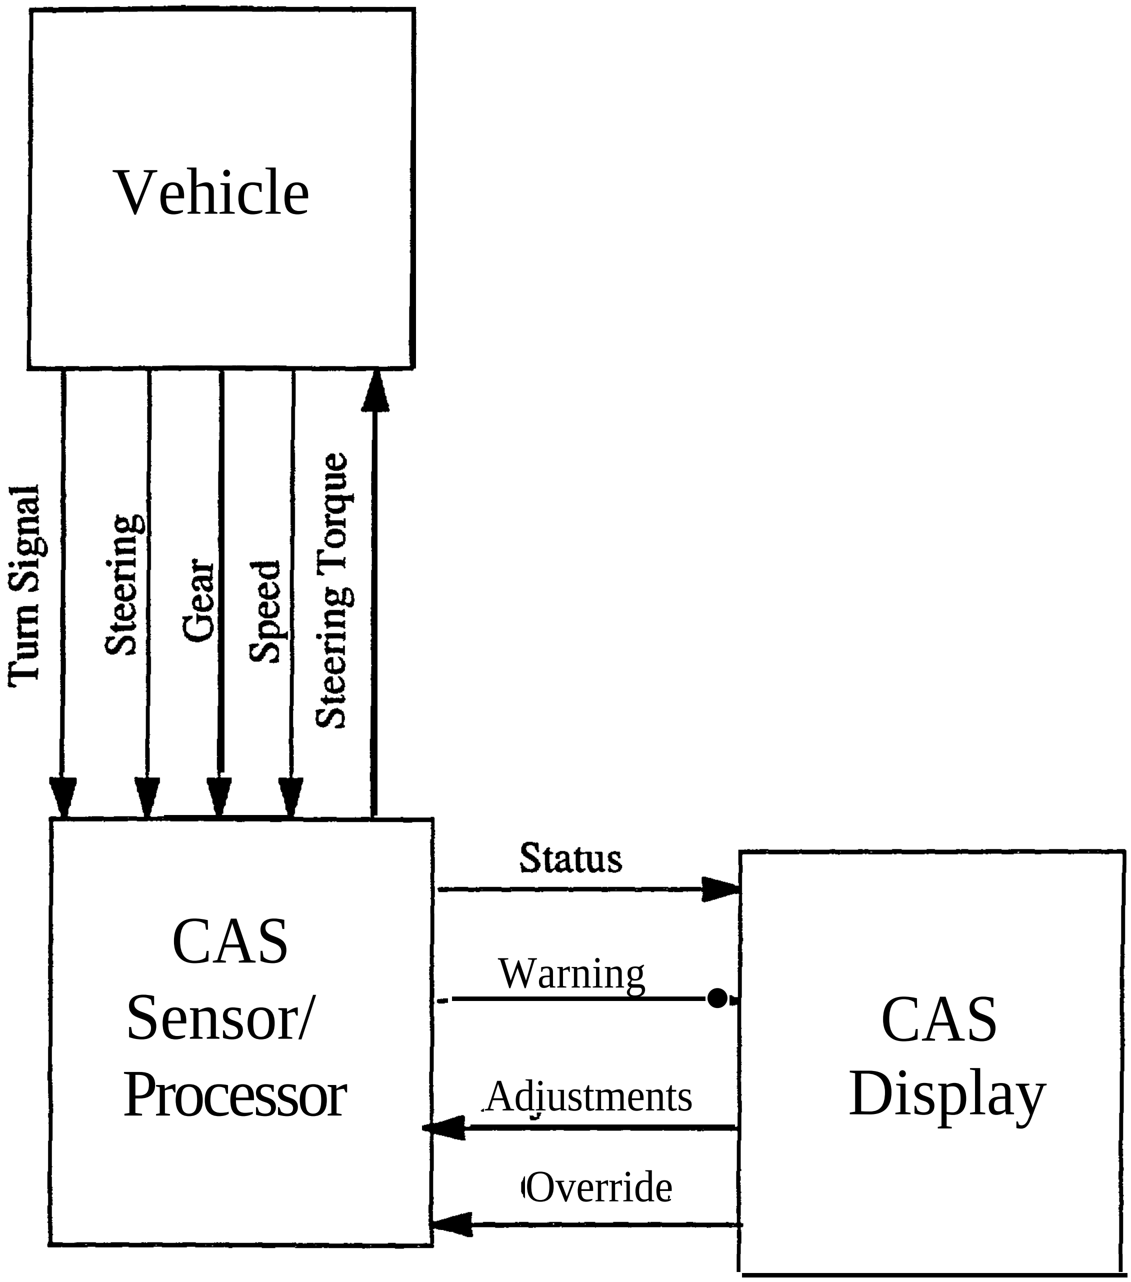

Figure 2.3.1 Interfaces between the vehicle, the CAS sensor/processor, and the CAS display.

vehicle. The gear information can also be utilized as part of the activation mechanism. For safety or interference reduction reasons it would make sense for the lane change/merge CAS to only be operating when the vehicle was in forward gear. Also, a minimum speed criterion might be useful for the same reasons. Finally, the steering torque feedback to the vehicle would be necessary if an “automatic” system were implemented. If it were determined that imminent danger would result from a lane change or merge maneuver, the steering of the vehicle could be tightened to inhibit or even prohibit that maneuver from occurring. Of course, a highly sophisticated CAS would need to exist for this extreme countermanding of the driver to be safely attempted.

The interface between the CAS sensor and processor subsystem, and the DVI subsystem would provide both status and warning data to the DVI. The status data would allow the DVI to alert the driver that the CAS is malfunctioning either by turning off a visual indicator or probably more appropriately by turning on a visual and/or an auditory alert. The key output of the processing algorithms will be a determination of the warning to be issued, if any, to the driver. This warning may have to convey a small number of severity levels and it may be useful to differentiate which zone (or at least which side) the threat was determined to be in.

The driver could also interface with the CAS through the DVI. The arrows going from the DVI to the sensor/processor represent this communication path. For example, there may be some form of sensitivity adjustments that would modify the operation of the sensor or more likely the performance of the algorithms. Warning zones may be modified based on the driving style of the vehicle operator, but only within prescribed limits. Also, an override may be provided to temporarily deactivate the warning for an ongoing but benign situation. These two paths have been shown for completeness, but it is not clear if either will be utilized.

The processor would analyze the sensor measured data in order to determine the level of warning to be issued if any. In addition, the processor would control periodic self-tests in order to determine the CAS status. A few bits of data would be periodically sent to the DVI to communicate the status and warning level to be displayed.

The requirements on the DVI include the ability to present clearly and unambiguously any warning determined by the processor subsystem. This impacts the location of the interface and the type and characteristics of the warning modality. Also, the status of the CAS must be presented to the driver. Other requirements are less certain. These include the ability to adjust some characteristics of the warning (such as, light or sound intensity) within limits, based on the ambient conditions in the vehicle and/or driver preference; the ability to adjust some characteristic of the warning algorithms to allow for different driving styles and driver reaction; and the ability to temporarily turn off the warning.

These last capabilities may not be confirmed as requirements until further human factor testing is performed on a variety of subjects.

2.4 Preliminary Performance Specifications

A summary of the derived preliminary performance specifications for the lane change and merging CAS that monitors both the area adjacent to the subject vehicle (SV) and fore and aft of the SV on either side out to one lane width is presented below. Figure 2.4.1 is a graphical representation of the coverage zones on the right side of the CAS-instrumented vehicle. This information has been extracted from our Task 4.4 Interim Report.

• Function:	Target Detection and Ranging; Longitudinal relative velocity; Driver alert
• Coverage:	1 lane (12 ft) to left or right in the transverse direction, depending on the lane change direction indicated; this coverage to extend in the longitudinal direction to 80 ft. (TBR) fore and aft of the SV; 1 - 10 ft in height; also for lanes that merge at angles up to 15^{0} (TBR)
• Size of Target:	Any vehicle allowed on public roadways (pedacycle to truck)
• Target Velocity:	Any'allowable (0 to 65 mph)
• Target Acceleration:	Any achievable (-g to +g)
• Number of Targets:	Presence detection of all targets, one or more per zone
• Platform (SV) Velocity:	Independent of SV speed
• Platform (SV) Acceleration:	Independent of SV acceleration
• Relative Velocity Range:	+/- 60 mph
• Measurement Latency:	Less than 0.5 s
• Measurement Accuracy:	**2** feet (range); 5 ft/s (TBR) (velocity)
• Performance:	Probability of Detection: > 99% (TBR) Probability of False Alarm: < 10^{-6} (TBR) Probability of Nuisance Alarm: < 10^{-3} (TBR)

- Interference including EMI/EMC: Shall not interfere with the operation of other in-board or out-board systems
- Duty Cycle: On-demand operation with TBD activation mechanism
- Driver Vehicle Interface: Headup display activated by turn signal or always-on visual signals in respective side view mirrors with audio warning tied to turn signals (TBR)

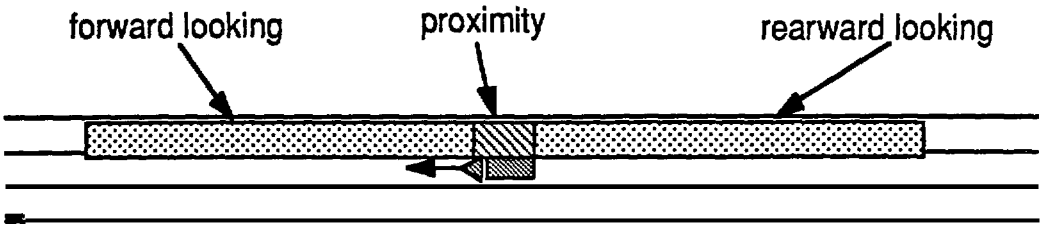

Figure 2.4.1 Coverage zones for a lane change CAS.

Of course, the detection of all targets in the shaded regions of Figure 2.4.1 would not be required. Only those targets that are closing at the appropriate rate with the instrumented vehicle (behind but going faster or ahead and going slower) need to be addressed. The closing rate deemed appropriate is a function of the reaction time of the driver, the speed at which the lane change/merge is accomplished, etc. Figure 2.4.2 graphically illustrates a representative warning decision region for the segment of the CAS system that is looking next to and behind the CAS-instrumented vehicle. The area denoted as “red” is the warning area. It includes a zone adjacent to the instrumented vehicle where a warning is issued regardless of the relative speed between the detected object and the vehicle. The warning zone extends to longer ranges for vehicles with positive relative velocities. The exact shape and boundaries of the warning zone area are matters of ongoing research. Many believe that some driver selectivity should be built into its determination.

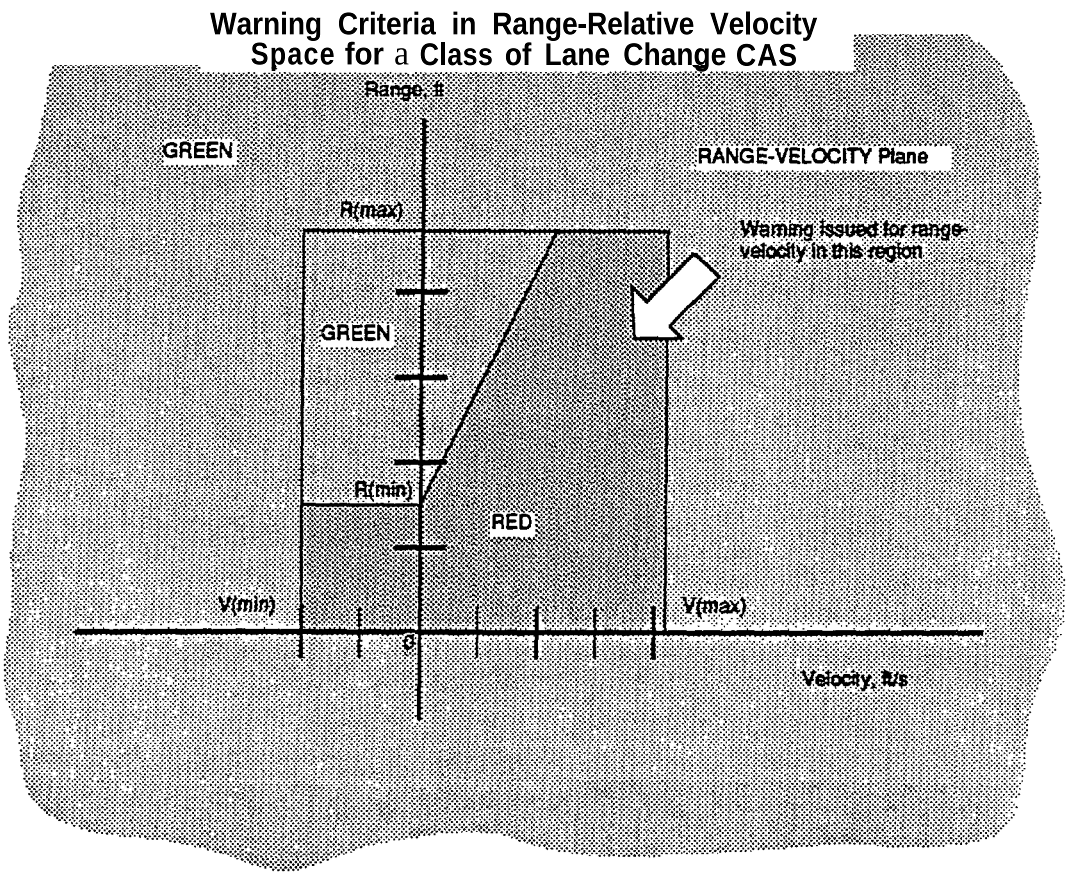

Figure 2.4.2 Representative decision region for a lane change CAS.

Additional specifications were addressed in our Task 4.4 Interim Report with regard to lane change/merge accident avoidance. These include ones for lane keeping to protect against drifting accidents, specifications for counter convergence to avoid accidents when two vehicles which are two lanes apart are simultaneously changing lanes into the one between them, and another for merging aid to warn drivers of the presence of other vehicles while merging into a lane from a non-parallel direction. In that specification, we suggested that the coverage zone adjacent to the instrumented vehicle should be expanded out to cover intersecting lanes at angles up to fifteen degrees (see Figure 2.4.3 for a representation). This expanded zone would open up the CAS to the possibility of many more nuisance alarms. In order to avoid this, determination of the trajectory of any target detected would be necessary in order to assess its potential threat during the lane change or merge maneuver. This needed trajectory computation, in turn, requires detailed knowledge of the targets position and velocity. The ability of the investigated sensor technologies to determine these characteristics will be addressed in Section 3.1.

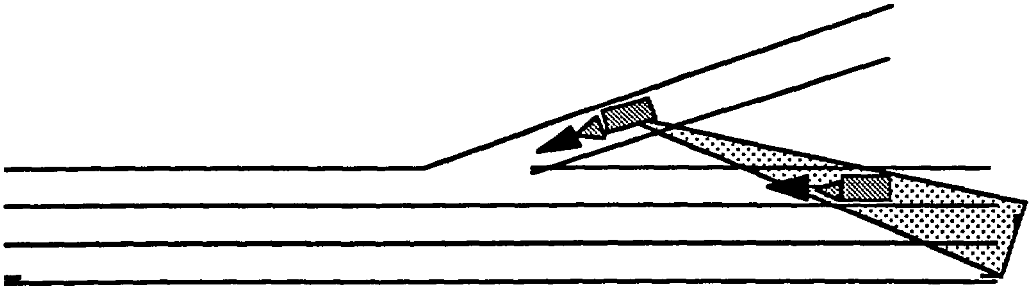

Figure 2.4.3 Coverage zone for a merge CAS.

3.0 Sensor Technologies

3.1 General Considerations

3.1.1 Active Versus Passive Sensing Systems

In order for a CAS to alert the driver of the presence of another object that may interfere with his or her planned lane change or merge maneuver, a sensor element must be present. Like our vision, this sensing must be done remotely. Our vision is passive in that it relies on reflected and emitted radiation that is detected by the eye. Passive systems are simpler in that they do not produce their own radiation source and hence there are no safety concerns nor interference problems. Their key drawback is that they require a very sophisticated processor (like the human brain) to interpret the incoming signals and derive the positions and directions of motion of the objects whose radiation is being received. One solution to that processing need is to involve the human's brain by presenting the passively received signal to the driver for interpretation. This approach has not been deemed acceptable because of its large potential for distracting the driver and/or overloading him or her with too much information. The other solution of providing an artificial brain (that is, a processor) to perform the difficult image interpretation tasks has not proven to be viable at this time nor does it appear to be solvable in the timeframe that we are considering.

This leads to evaluate active sensing systems, that is, ones which produce their own energy source and detect reflected energy from objects in the sensor's field of view. Field of view usually refers to the area in space that the sensor can nominally detect objects in when it is at one fixed position, and field of regard is defined as the total area investigated by the sensor as it is scanned or moved. Conceptually the simplest form of an active sensor is one which emits a short burst of energy and then detects its return after its interaction with objects in the field of view. Since the speed of propagation of the radiation is known, then by measuring the time between the start of the emitted pulse and its return, the distance to the object is readily determined. As can be easily understood, the accuracy of this measurement is proportional to the duration of the pulse, and hence short pulses are necessary for accurate distance measurements. The distance to the target is referred to as the range and the measurement is called ranging. Conceptually, the simplest way to measure the speed of the target is to measure successive ranges and then deduce the speed from their differences. Another approach is to rely on the Doppler shift which occurs when a wave is emitted or reflected from a moving object. The change in frequency due to the motion is proportional to the velocity along the line of sight from the receiver of the radiation to the emitter/reflector. Note that both the range and speed measured in these ways are the “radial” component.

3.1.2 Technology Choices

Common choices for emitted energy types for remote sensing include microwave and millimeter wave for radars, visible and infrared for lidars, and acoustic and ultrasonic for acoustic sensors. No known acoustic sensor can provide the range necessary for the requirements established in Task 4.4 for a lane change sensor, and so we will not include that category of sensor in our discussions. Applications of acoustic sensors for shorter range systems, such as for a backing CAS are reasonable and it is expected that further development in that area will be forthcoming.

A note about terminology. Radar stands for radio detection and ranging. We have chosen to use the acronym lidar (light detection and ganging) as opposed to ladar laser detection and ranging) or the cumbersome laser radar because of its consistency with the original meaning of the term radar.

Radar was the first active sensor system developed because of the availability of radio wave sources. The early radars operated in the microwave region with frequencies well below 1 GHz. As technology improved, radars began to appear that operated at frequencies up to 10 GHz. The increase in frequency is important because of radar’s major limitation which is beamwidth. A system that receives radiation (typically an antenna for radio waves and a lens for light or infrared) can differentiate the direction of that radiation over an angle (angular resolution) proportional to the wavelength of the radiation and inversely proportional to the aperture size. Because radio waves have wavelengths that are orders of magnitude larger than visible light (typically on the order of centimeters for microwave and millimeters for millimeter waves versus microns for light and infrared) they require much larger receiver apertures. A factor of ten increase in frequency (which corresponds to a factor of ten decrease in wavelength) allows for the reduction in diameter of the antenna by a factor of ten (and in area by a factor of 100) to maintain the same angular resolution.

The utilization of light and infrared radiation for ranging purposes (lidar) was a direct consequence of the development of lasers. All lidars utilize a laser as the source of the radiation to be propagated and detected. The key advantage to a lidar over a radar is its very high angular resolution stemming from the much shorter wavelengths employed. The key advantage of a radar is its much better penetration capability through limited visibility conditions (fog, dust, smoke, light rain) than lidar due to its much longer wavelength. Of course having a much tighter coverage zone in angle is not always an advantage, such as, when large areas need to be investigated. Also, propagation issues are not as important in very short range applications.

3.1.3 Detection

Much of the discussion in Section 3.2 will be couched in the language of radar since the measurements techniques presented were first developed for radar sensors. Four excellent books on the subject of radar that cover these topics are contained in the references (6-9). In order to detect the presence of a target and measure its position and/or speed, sufficient energy must be scattered from it and received by the sensing system. In other words, the energy returned due to the presence of a target must be higher than the noise in the receiving system by an amount large enough to be discriminated from the noise. This noise comes from a variety of sources, such as the ambient background, the internal electronics, etc. The standard measure utilized to determine if detection will occur is the signal to noise ratio. If the target return (the signal) is larger than the ambient energy (the noise) then a thresholding of the measured power can be employed to discriminate against the noise. This threshold must be set high enough to eliminate most of the noise but low enough to detect most of the targets. This contradictory requirement is what makes the detection problem so difficult, especially when very high probabilities of detection are required. The detection probability is defined as the number of targets detected divided by the number of targets present. Of course, the types of targets that must be detected need to be specified, as well as the area over which the detection must occur. The false alarm rate is the frequency of time that noise is reported as a target detection.

In addition to false alarms, another kind of false detection may occur. This happens when the detection threshold is crossed but the detected energy is not from an appropriate target. Here appropriate means that the target is not in the class to be detected or it is not in the region where detections should occur. That region can be physical space or it can be phase space, where one or more coordinates represent position, and one or more represent speed. These nuisance alarms can occur because the emitted energy from the system is not focused enough over the region of interest, the determination of the range and/or velocity is not accurate enough, or the radiation interaction with the target cannot discriminate between valid and invalid targets. The susceptibility of the different sensor technologies to false and nuisance alarms, and the robustness of the processing approaches to discriminating between true targets and spurious ones (either because of noise or inappropriate returns) will be evaluated in the coming sections.

The received signal consists of radiation scattered by all objects "illuminated" by the radar including cars, trucks, the road, trees, signs, the atmosphere, precipitation, etc., as well as, radiation generated by the electronics itself (noise). The returns from the uninteresting background is termed clutter. The radar illumination pattern is determined by the antenna and includes the area of the main beam and, to a lesser extent, the sidelobe areas.

The usual detection process involves the measurement of the maximum amplitude of the signal (the envelope). For a normally distributed signal, the envelope will display a Raleigh distribution, which has a simple one parameter exponential form. Thus, if a signal contains only noise or clutter returns, the statistics of the noise plus clutter can be determined by calculating an average. Given this distribution, it is simple to set a threshold such that the clutter and noise returns will only be above that threshold a predetermined amount of times. This is the basis of constant false alarm rate (cfar) processing. A false alarm is defined as a case when the clutter plus noise return is greater than the threshold. Clearly, by setting the threshold high enough, the false alarm rate can be kept as small as desired, but, as will be shown, target detection is penalized. Mathematical details about the detection process are included in the appendix.

When an object's return is present in the range profile, its signal to noise ratio (snr) and signal to clutter ratio (scr) will determine how far it rises up out of the background. Given a desired probability of false alarm (pfa), the threshold is determined. The probability that the object is detected (pd) is then set by the snr and scr. A detected object is then termed a target. A low pfa requires a large threshold, which leads to missed objects (low pd), and, conversely, a large pd demands a lower threshold which gives rise to a higher pfa.

The simplest way to determine the statistics of the background (noise plus clutter) is to average all the range cell returns. This is the cell average cfar (ca-cfar) technique. In practice, a single range cell is selected, the signals from the others are averaged, and the return of the cell under consideration is compared to the average times a factor set by the desired pfa. This approach is adequate if the background is relatively constant and if there are only a few targets. A varying background can cause the threshold to be too low in areas of high background and too high in areas of lower background. The presence of many targets can corrupt the average and lead to a high threshold. Many techniques have been developed to alleviate these shortcomings of simple ca-cfar. They include greater of cfar (go-cfar) and ordered statistics cfar (os-cfar).

The detection of any possible threat to the CAS-equipped vehicle is critical. This drives the algorithms to very high pds. However, this will also lead to relatively large pfas which may cause the driver to ignore the CAS's warnings. There are techniques for increasing the pd without the burden of greatly increased pfa which utilize dwell-to-dwell comparisons. A dwell here is defined as the time over which the data necessary to form a range profile is accumulated. Non-coherent integration is one such procedure. Basically, the measured signals from a number of dwells are added before the ca-cfar step, thus reducing the random components of the background relative to the more stable object returns. Also, an "m-out-of-n" approach could be utilized in which say 4 out of 10 threshold crossings must occur in consecutive dwells before a cell is identified as containing a target.

3.1.4 Ranging

As mentioned above, the simplest method to measure the range to a target is to bounce radiation off of it and measure the time it takes for the roundtrip. Since there is always uncertainties about which part of the emitted pulse is being scattered off of the target, the duration of the pulse sets the range resolution of the sensor. The range resolution is given by the propagating speed, the speed of light for both radar and lidar, times the pulse duration time divided by two because of the roundtrip. For example, to get a range resolution of one foot, a pulsewidth of two nanoseconds is required. In order to avoid ambiguities arising because returns from one pulse are being received after at least one more pulse has been transmitted, a maximum prf (pulse repetition frequency) is set to allow for the return of radiation from the maximum range ever expected before the next pulse is generated. Because of the shortness of the pulses, and the requirement to wait between pulses, it is sometimes difficult to get enough energy on a target (to have enough signal to noise ratio) to effectively detect it.

Two solutions to this problem have arisen. One involves simply adding (integrating) the returns from a number of pulses before performing the detection operation. Coherent integration enhances the signal to noise ratio to N times that for a single pulse, where N is the number of pulses added. For non-coherent integration the increase is reduced to the squareroot of N. Another approach to enhancing the energy returned from a target is to "code" a longer pulse and then "compress" the data to achieve high resolution.

Two forms of coding that are regularly utilized in radar will now be discussed. The first involves varying some feature of the waveform in discrete steps during a single long pulse. This coding can involve the phase or amplitude, for example. In this way, even though the time duration of the pulse is long, its return from the target can be correlated with a segment of the outgoing waveform by identifying the coding sequence and thus time accuracy to a small fraction of the waveform duration can be achieved.

The other form of coding is quite regularly used for automotive applications. The reason for this will be discussed later. It employs the linear FM (frequency modulation) waveform also called the chirp waveform. During the elongated pulse, the carrier frequency* is linearly changed from some value to another one. The difference between the starring frequency and the final one is the bandwidth. It can be shown that the range resolution of this waveform is independent of the pulse duration and equal to the speed of light divided by twice the bandwidth. What makes this waveform appealing for many applications is the low peak power required because of the relatively long duration of the chirp (typically tens of microseconds to milliseconds), and the ease in which the processing can be performed to derive the range. This will be discussed in the processing sections of this report. A variant of this waveform is the stepped frequency one where the frequency is changed in discrete steps. Again, the total bandwidth spanned determines the resolution.

3.1.5 Speed Measurement

Two standard ways of determining speed with a ranging device involve determining the rate of change of the range and Doppler processing. The first is straightforward in that it utilizes whatever waveform is employed to measure the range. The only addition to the sensor system is additional processing in which sequential range measurements are differenced to derive the range rate. The accuracy of this computation is proportional to the accuracy of the range measurement and inversely proportional to the time interval employed. Although this approach to speed determination is the simplest, it can lead to large uncertainties in the velocity determination, especially when rapid updates are required.

The other approach is to directly measure the frequency shift that occurs when a source of energy is moving relative to a receiver. This is the Doppler effect, the classic example of which is the shift in a train whistle as a train approaches and passes an observer at a station. If the range between the source of energy and receiver is changing, then the wavelength and frequency of that received energy is shifted, the magnitude of the shift being proportional to the relative velocity between them. For a ranging device, the energy source is the reflected energy from the target and the relative motion can be caused by either the receiver or target moving, or both. The simplest way to measure this frequency shift is to utilize a waveform consisting of a pure tone (frequency) and then employ a method to determine the difference in constant frequency between the received and transmitted energy. This is incompatible with the ranging process since a constant frequency waveform intrinsically contains no timing information.

Books have been written on how one can measure both range and Doppler with ranging devices (for example, see references (6-9)). A few approaches will be summarized here, One straightforward way is to employ two interleaved waveforms, one that measures range and the other which measures speed. There are several disadvantages to this approach. Clearly, it’s complicated. Secondly, the signal to noise for each type of measurement is reduced since timesharing of waveforms is occurring. Perhaps the most important drawback for automotive applications is the ambiguity involved when multiple targets are present. One must correlate the ranges measured during one part of the waveform with the velocities measured during the other. Again, a vast literature exists discussing this topic (again see references (6-9)).

Another method is to utilize coherent pulses and to integrate the returns from many pulses to determine the Doppler while the pulses can be utilized to determine the range directly. Coherent is used here to mean that the phase relationship between the various pulses is controlled as if they were all cut out of a single wavetrain. The requirement of coherence causes the hardware of the ranging device to be complicated and can add greatly to the expense of the system.

Finally, the linear FM waveform mentioned earlier can be employed. Here again there is an ambiguity involved between the range and Doppler measurement. Multiple chirps and chirps interleaved with constant frequency periods have been employed to alleviate this issue.

3.1.6 Bearing Measurement

For some more sophisticated forms of a CAS, detailed knowledge of the trajectory of the detected object may be necessary. For example, drifting accidents could be prevented if the motion of the one vehicle across the lane boundary could be detected. Also, for merging accident avoidance, where wider-angle zones of coverage are necessary, the prediction of the intersection of your trajectory with the detected vehicle’s path is the driving criteria for issuing a warning.

Another measurement that can be made by a ranging sensor that will enhance the ability to determine the target’s trajectory is the bearing. Bearing is defined as the direction angle of the target from some pre-defined standard direction. The accuracy of this measurement is about the beamwidth for fixed antenna systems, but it can be greatly enhanced for scanning ones. By measuring the relative strength of the return from the same target at two slightly different angular pointings, precise knowledge of the direction of the target from the instrumented vehicle can be ascertained. This can be done by sequential measurements (for example, with a conical or raster scan, or with sequential lobing) or by simultaneous ones (monopulse, where several different measurements are made simultaneously by utilizing sections of the antenna).

A determination of the time rate of change of range and bearing leads to complete knowledge of the motion of the target. Thus the parallel and perpendicular velocity can be deduced and utilized as part of the threat assessment computation. Lidars, with their intrinsically smaller beamwidths, can make more precise bearing determinations.

3.2 Radar

3.2.1 Basic Approaches

As discussed above, there are several approaches to measuring both range and velocity with a ranging sensor. The ones most commonly implemented in automotive applications for a radar are linear FM and pulsed waveforms.

As will be discussed below in Section 3.3.2, the use of millimeter wave monolithic integrated circuits (MMIC) is enabling the manufacture of smaller more reliable, and lower cost RF circuits. Integration is now to the stage that an entire transceiver (combination transmitter and receiver) can be produced on a single chip. MMIC devices tend to operate best in higher duty cycle, lower peak power modes. They are most compatible with linear FM waveforms for this reason. In addition, the lower peak power required is beneficial from a safety viewpoint. Because of these factors, there has been a tendency to develop a number of automotive radar systems with that approach. Lower frequency applications can use more conventional discrete component devices and therefore are not as strongly inclined towards the linear FM waveform.

Anyone who has looked under the hood of a new car knows how cramped for space they are. In addition, miles of cabling are woven through every available space throughout the body of the vehicle. Any CAS system added to a car must compete for the very limited space available. In addition, a CAS (or at least the sensing part of it) must be positioned to “see” the areas to be monitored. For the lane change/merge application this means visibility at both sides of the vehicle. Potential mounting points include the side view mirrors, the sides of the vehicle, or the rear area.

Since the sensing part of the CAS must be able to “see” the areas around the car, it must be mounted on or near the surface of the vehicle. Hence it must also meet the additional constraints of not interfering with the styling, and equally important, being compatible with the external environment that a car is exposed to. The best way to make the CAS compatible with all of these constraints is to make it as small as possible. This is the great advantage of a lidar which tends to be smaller than any radar devices. Remember, that the aperture size is proportional to the wavelength and that laser wavelengths are many orders of magnitude smaller than those associated with radar. The easiest way to make a radar small is to utilize as short a wavelength (as high a frequency) as possible, This keeps the aperture smaller, allows for smaller components, and provides for additional bandwidth which provides for better range resolution for linear FM waveform approaches. A simple rule of thumb is that the percentage bandwidth available at a given frequency is constant and therefore higher frequencies provide for larger bandwidths.

The choice of frequencies for automotive radar applications has varied from about 10 GHz up to 94 GHz. Safety is not much of an issue for frequency selection since the energy density requirements for RF are the same over that entire frequency range. What drives the frequency choice is the availability of components, the size of the radar (especially the aperture), and the atmosphere propagation characteristics.

For longer range radar applications, designers had stayed below about 18 GHz to take advantage of existing hardware and to avoid the increased atmospheric absorption which in a general sense increases with increasing frequency. There are so called “atmospheric windows” for RF transmission which can be found around 35 and 94 GHz. These windows come about because of peaks in the absorption associated with specific atmospheric constituents. The largest occurs at 60 GHz and is due to a resonance in the 02 molecule. Because of this large absorption, the band around 60 GHz has been reserved for space-to-space restricted communications and to very local area communications systems on the ground. After the large peak at 60 GHz, the absorption decreases for a while until the trend with frequency again takes over. This leads to a dip around 94 GHz. This frequency band has been exploited by the military, especially in seeker applications. In order to guide a missile to a target autonomously, it must be able to “see” the target. This requires mounting a sensor in its nose cone. In order to get reasonable angular resolution with the restricted aperture, the tendency has been to utilize the highest frequency radar available. At the present time that is a 94 GHz one.

In Japan, the Government has designated the 60 GHz band for automotive applications. This was done to minimize interference because of the large attenuation of the signal through the atmosphere. Also, there is very little utilization of that band by other applications. In Europe, the 77 GHz region has been chosen. This is a compromise between the desire to be in the 94 GHz window and the wish to have a band dedicated to only automotive applications. The attenuation at 77 GHz is not much worse than at 94 GHz In the United States, a variety of bands have been requested for licensing from the FCC (Federal Communications Commission). The decision to exclusively allocate a number of bands is being reviewed presently (l0). These bands range from about 40 GHz to 153 GHz. (There is another atmospheric window around 140 GHz.) So far, the band at 46.7 to 46.9 GHz and the one at 76.0 to 77.0 GHz are being proposed for exclusive use for automotive radar systems (FCC 95-499 dated 12/15/95 and reported in the Federal Register of 4/2/96).

A 77 GHz linear FM radar will now be utilized as a representative case. In order to obtain a snr of 15 dB (necessary for the required detection probability) at 80 ft, 5 mW of power need to be transmitted. This calculation has been made assuming a 1 m2 (0 dBsm) target cross section which is conservative. Most cars have cross sections between 5 and 10 dBsm. A system noise figure of 18 dB has been employed, which is again conservative. The antenna gain has been assumed to be 25 dB. This is consistent with a beam size which is one lane width (12 ft) at the maximum range of 80 ft. Finally, the chirp duration employed is 33 us.

3.2.2 Technologies

The use of millimeter-wave (MMW) radar for automotive collision avoidance is not new. Many companies, institutions and laboratories have experimented with various MMW pulsed and/or FM-CW radar systems. Table 3.2.1 summarizes some of these efforts

1980	TOYOTA-Fujitsu	50 GHz FM-CW
1982	NISSAN	60 GHz Pulse-FM, NRD- Guide
1988	Phillips	94 GHz FM-CW, Low cost Hybrid
1990	Univ. of de Lille	94 GHz Quasi-optical
	SMA	38.5 GHz, Pulsed, Noncoherent
1991	Phillips	77 GHz FM-CW, Low cost Hybrid
1992	GEC-Plessey	77 GHz FM-CW Quasi-optical
	TRW	94 GHz FM-CW, single chip
	Fujitsu	60 GHz FM-CW
	Univ. of Munchen	61 GHz PN coding
	TEMEC/DASA	77 GHz MMIC coherent pulse
	Lucas Ltd.	77 GHz FM-CW Hybrid
1993	Millitech	76.5 GHz Pulsed/FM-CW
	DASA	77 GHz FM-CW, Low cost Hybrid
1994	Celsius Tech	77 GHz FM-CW hybrid
	Phillips	77 GHz, frequency scanning
	Raytheon	77 GHz, electronic scanning
	HIT	77 GHz FM-CW

Table 3.2.1 A representative summary of automotive radar approaches.

For automotive radar applications, a very popular radar waveform to use is the Triangular Frequency-Modulation Continuous Waveform (FM-CW). For this waveform, the range is determined by the instantaneous frequency difference between the transmitted and returned signal, and the velocity (or Doppler shift) information can be calculated by knowing the frequency difference between the up-sweep and down-sweep. The triangular FM waveform also allows unambiguous resolution of multiple targets. If two targets are present, they produce two echo signals at the radar. A spectrum analysis of the up-sweep frequency differences and the down-sweep frequency differences gives the individual targets in different frequency bins. To avoid the frequencies being “scrambled”, a third waveform segment is added which has no modulation.

Another type of radar also has been used. It illuminates the targets using short bursts of energy and then listens for echoes with the transmitter silent. This type is known as a pulsed radar. Time-of-flight information is utilized to deduce range and the coherent processing of a number of pulses can yield velocity. The pulsed radar requires higher transmit output power. This raises the question of whether pulsed radar output power is compliant with FCC regulatory requirements in regard to unlicensed radar operation output power restrictions.

The most significant technology advancement in the past few years affecting millimeter-wave radar is the revolutionary development of millimeter-wave monolithic integrated circuit (MMIC) technology. The ability to monolithically integrate several MMW functions into a single chip and to produce them in a high volume semiconductor process holds good prospect for low cost production of millimeter-wave radar modules which are the essential element for the creation of an economically feasible automotive radar market. Efforts are currently underway in several U.S. and international R&D firms to develop millimeter-wave MMIC for automotive applications.

Most of the millimeter-wave MMIC development work to-date is based on pseudomorphic (PM) high electron mobility transistor (HEMT) device technology. The HEMT device has demonstrated high gain, low noise and adequate output power capabilities and is particularly suited for millimeter-wave transmit and receive applications. Using this technology, a single chip W-band FM-CW transceiver has been demonstrated (l 1). It includes four circuit elements: a 94 GHz voltage controlled oscillator (VCO), a 2-stage output amplifier, a 3-stage low noise amplifier (LNA), and a mixer. The chip exhibited >10 mW of output power and a 400 MHz frequency tuning range at 94 GHz. The chip measured 6.9 mm x 3.9 mm. Redesign and layout of the chip based on new processes can reduce its size to about half of that.

Recent advancement in selective molecular beam epitaxy (MBE) technology has demonstrated the capability of monolithic integration of HEMT and HBT (heterojunction bipolar transistor) devices on the same substrate. By combining the high gain and low noise figure of HEMT devices and the low l/f noise and high linearity of HBT devices, optimum RF performance can be realized from a single integrated circuit (IC). The HBT is grown first by MBE, patterned with silicon nitride, and etched to form HBT islands. The pseudomorphic InGaAs-GaAs-AlGaAs HEMT material is then deposited. A comparison of the performance of a HEMT low noise amplifier fabricated by the HEMT-only process and one fabricated using the merged HEMT-HBT process (12) has shown essentially identical gain and noise figure performance. Work is currently underway to design a W-band FM-CW transceiver IC based on this HEMT-HBT IC technology. Significant improvements in VCO phase noise and tuning linearity are expected from this emerging technology.

There has also been some interest in exploring higher millimeter-wave frequencies for automotive radar applications. The advantage of going higher in frequency is the ability to achieve better spatial resolution with the same antenna aperture. The next propagation window above the 94 GHz one is at 140 GHz. There has been excellent progress made in developing 140 GHz MMIC components. Recently demonstrated is a 140 GHz monolithic two-stage low noise amplifier using 0.1 um InAlAs/InGaAs/InP pseudomorphic HEMT technology (13). This MMIC amplifier has a measured gain of 9 dB and a noise figure of 5 dB at 140 GHz. Developing a single chip FM-CW 140 GHz transceiver is feasible by integrating an amplifier like this with a 70 GHz VCO, a 70-to-140 GHz frequency multiplier and a 140 GHz harmonic mixer,

For automotive radar applications, besides the typical range, speed and angular position information, it may be important to be able to track multiple targets under different conditions: on crowded highways, in the presence of adjacent cars and around curves. To meet these requirements, a scanning antenna system may be required. There are several methods to accomplish the scanning antenna function:

- Electronically phase scanned array antenna - by using a phased array to produce a narrow beam and rapidly point the beam in specified directions allows tracking many targets simultaneously and with great precision. The key component in the phased array is the electronic phase shifter. Phase shifters are commonly realized with either ferrite phase shifter technology or p-i-n diode phase shifter technology. The advantages of electronic scanning are fast scanning, and no moving mechanical parts. The disadvantages are high loss, sensitivity to temperature variation, and potentially high manufacturing costs.
- Switched beam scanning antenna - the scanning of this type antenna is done by sequentially switching the different beams. Beam switching antennas have been realized with both quasi-optical techniques and waveguide switchtree structures. This approach also suffers high loss especially as the number of switched beams increases.
- Mechanical scanning antenna - this class of antenna includes gimbaled antennas, cassegrain type antennas with a scanned main reflector, rotating slot arrays, and other

mechanically-driven scanning approaches. The advantages of this type of antenna is low loss and mature technology. However, they suffer from slow scanning speed and have reliability concerns due to the moving mechanical parts.

Electronically phase scanned array antennas have been successfully demonstrated at 77 GHz. Both switched beam and mechanical scanning antennas have been demonstrated up to 94 GHz.

3.2.3 Safety Issues

Current ANSI standards restrict exposure to RF energy to levels below 10 mW/cm^2 over the entire range from 15 to 300 GHz. This restriction is based on acceptable levels of total heat input to the body. The FCC is proposing (l0) to adapt a more stringent requirement, namely the one set by the National Council on Radiation Protection (NCRP). Their maximum safe power density is given as 1 mW/cm^2 over the entire frequency range of 1.5 to 100 GHz. For a detailed discussion of safety issues regarding exposure to both microwave and infrared electromagnetic radiation, see reference (14) and references therein.

For the example linear FM system, only a peak power of 5 mW is required to achieve the desired snr. Thus for an aperture greater than 5 cm2 in area, the safety restriction is met. This implies a circular aperture with a radius of 1.3 cm. This is easily accommodated in any reasonable design. On the other hand, a pulsed system would utilize significantly larger peak powers to generate enough energy on target. This raises the issue of safety, at least in the vicinity of the transmitting aperture. In either case, some care must be given to restrict the removal of the transmitting aperture to expose the RF energy feed which typically might be a waveguide. The energy density at the opening of the waveguide would be quite high since it is very small especially at millimeter wavelengths. It is recommended that some interlock control be provided that would preclude transmitting when the waveguide opening is exposed.

3.2.4 Interference Issues

There are various interference issues associated with a CAS. These involve interference between the same types of CAS on different vehicles, different types of CAS on the same or different vehicles, interference between the CAS and some other system on the instrumented vehicle, or interference from some external source of radiation and the CAS. There are standards involving any electrical equipment that are integrated into vehicles (for example, see references (15-16)). These cover the amount of radiation that can be emitted by the equipment and the ambient radiation levels that it must be designed to operate in. Clearly these same standards will hold for a CAS. However, unlike other system, the CAS will be a source of energy that is directed toward other vehicles on the roadway. The potential for this radiation to couple to another vehicle's CAS system or other CAS on the instrumented vehicle is the driving interference issue that must be confronted.

It has long been recognized that interference will be a major concern when a large number of CAS are deployed. The critical reason for this is a simple one. The energy radiated from a source has an intensity that falls off like one over the distance squared from the source. When it is reflected from the target, it again has the same decrease with distance from the target. This leads to the well-known "one over range to the fourth" dependence in radar analysis. Besides this factor which greatly reduces the amount of power received by a radar, there is the cross section of the target itself which also reduces the power received. Compare this with the energy received from a different radar. If the other radar is pointing towards the receiver of the first one, then only the reduction from the one over the distance squared term occurs. This implies two disturbing facts. One is that the energy from another radar can be much larger than the return from the receiver's own transmitter. Second, the range over which the energy is comparable is much larger than the maximum operating range of the CAS itself.

A variety of ways to mitigate this potential interference issue have been described in the literature. These include polarization adjustments, coding approaches, frequency diversity, waveform choices, and cooperative systems. All of these will be discussed in order.

One of the first approaches towards mitigating against interference from approaching vehicles is the use of linear polarization which is rotated by forty five degrees from the vertical. Polarization refers to the direction of the electric field in the electromagnetic radiation. Linear polarization in which all the electric fields are aligned is the simplest form. If all the antennas were aligned at forty five degrees from the vertical, then a very large fraction of the reflected energy will have the same polarization and be readily detected. Whereas, radiation from an approaching vehicle would be orthogonally polarized and would not couple very well to the other vehicle's antenna. This is a good solution for oncoming interference, but does nothing to reduce interference from reflected energy from other vehicles' radars which are traveling in the same direction.

As mentioned in Section 3.2.4, coding techniques have been utilized in radars to increase pulse duration and still achieve high range resolution. By utilizing a variety of codes and deploying them randomly (as is done for example with garage door openers), the risk of interference is greatly reduced. Codes exist which are highly orthogonal, and hence the energy received from another system would not be interpreted as a valid response from a target. Also, if another radar's chirp is received in a linear FM system, it would most likely not be mixed down to an acceptable IF and thus not produce a spurious detection. Note that this approach is intrinsically different from the polarization-dependent one. In the former, the interfering energy is not detected by the receiver, whereas in the latter it is received but discriminated against. For interference reduction techniques in the second category, there is an additional concern. Even though the interfering energy source is not interpreted as a target, it still couples into the receiver, and thus if it is large enough it can damage or at least saturate the receiver. Although damage is unlikely, saturation is not. When a receiver is saturated, then it cannot perform correctly for a period of time.

Frequency diversity is a natural way to avoid some interference, since receivers can be designed to accept energy only in a prescribed band. Thus, if a variety of widely separated frequency bands were available, interference could be greatly reduced. There is a problem with this since, as we have discussed, there is only a restricted set of frequency bands that will be allocated to automotive applications. Thus the number of discrete wave bands available to the designer will be limited, especially for the wider bandwidth approaches like linear FM modulation.

Waveform choices play a key role in the design of any radar-based system. They also impact the amount of interference that can be expected. As previously discussed, linear FM and other coding can be utilized to avoid some interference. In addition, a low duty cycle (defined as the fraction of the time that the system is actually radiating) pulsed radar would avoid a significant amount of interference just by being time-gated. Again, there is still a susceptibility to saturation even when the radar is not expecting an energy return.

The last approach is the least practical and most costly. It requires every vehicle on the roadway to have one or more transponders on them. A transponder receives radiation and then t-e-radiates it. In order to avoid some forms of interference, it would be designed to receive at one frequency band and re-radiate at a completely different one. In this way, interference from vehicles going in the opposite direction would be eliminated since their radiation would be completely out of band of the expected radiation from the transponder. Of course, this doesn't eliminate the possibility of one vehicle receiving transponder output stimulated by another vehicle's transmission. Because, this approach requires all vehicles (including ones which are not instrumented with a CAS) to be equipped with a transponder, its implementation is highly unlikely.

It is anticipated that the first approach to interference avoidanoe will be made through the judicious choice of waveform. Low duty cycle systems or ones that utilize coded waveforms are intrinsically less susceptible to interference. The use of polarization alone to avoid interference will not work since the geometry of the roadway is so complicated with many possible scatterers and many emitters moving in a variety of directions. Electromagnetic spectrum allocation is limited and the use of frequency variation is probably not feasible when systems become numerous. Finally, cooperative systems will not be implemented for many years, if ever, because of the necessity for all vehicles to be consistently instrumented.

If the choice of straightforward waveforms does not appropriately limit the risk of interference as these CAS systems become widely deployed, then additional waveform coding will be necessary. Careful monitoring of the performance of these systems will have to be maintained to ascertain if the interference risk is increasing to significant levels as deployment increases.

3.3 Lidar

3.3.1 Basic Approaches

The basic approach that has been employed in automotive laser radar is to use a laser source such as a semiconductor laser to provide a range measurement. For a general discussion of the use of lasers in remote sensing, see reference (17). Range is determined from laser power scattered from a target and detected by an optical power detector, so-called direct detection. A number of options exist for determining range. The most common involves measuring time of flight for a short < 100 ns laser pulse. Another method involves continuous sinusoidal amplitude modulation (AM) of the laser. Here range is determined from the relative phase between the outgoing and return beams. A third method involves triangulation. In this case to obtain adequate precision the transmit and receive apertures are separated by some distance.

As indicated in the previous paragraphs, range and velocity measurements are necessary to successfully analyze collision threats while minimizing nuisance alarms. Coherent lidars have been developed which can measure velocity directly through Doppler shifts, but because of their greater cost and complexity these technologies have not been favored for automotive applications. This leaves lidars based on direct detection in which only range is measured. Relative velocity information is then derived from a sequence of range measurements.

A large number of lidar systems have been built by various experimenters over the years. Many automotive companies have investigated lidar for collision avoidance. Most employ a semiconductor laser at around 900 nm in pulsed mode, such as the European Prometheus system (18) and several Japanese systems (19,20). Pulsed systems tend to have a limited minimum range, even though they have separate transmit and receive apertures. Factors which limit the minimum range are aerosol and obscurant backscatter, and accuracy of the system time of flight measurement for short times periods. Generally, the minimum range cannot be less than the system range resolution. Often, clock rate is a limiting factor when clocks are triggered on or off by simple pulse edge detection schemes, and the clock period is greater than or comparable to the laser pulse width. Pulse systems have generally not been reported to range accurately at distances which correspond to flight times less than the laser pulse width, because for edge triggered clocking, the pulse width usually limits the range resolution. We don't believe this represents a fundamental technological limit. Lidars which employ digital sampling and centroiding of pulses should be able to achieve range resolutions of a fraction of the lidar temporal pulse width, as has been shown recently for a space-based altimeter. This range resolution will then establish the minimum range.

Automotive applications of amplitude modulation are less often found, although the technique has found commercial use for optical inspection. Workers in China reported results with an AM-based automotive lidar (21). Generally, AM systems work well at close range, although there is still a minimum range based on the minimum resolvable phase shift that the signal processing electronics can detect for a given snr. Generally, AM systems are limited in maximum range by the range ambiguity of the modulation waveform:

$$R_{max} = \frac{c}{2f_{AM}} \quad \text{where } f_{AM} \text{ is the modulation frequency.}$$

In the AM system, considerations of range resolution and maximum range tend to oppose one another. Higher range precision can be achieved using higher modulation frequencies, but this decreases the maximum range.

Triangulation systems find almost no use in automotive lidar. There is mention of using a triangulation scheme for the Prometheus system in order to measure short ranges, since the system'was limited to a minimum range of 10 m at the time. This was due to limitations of the analog receiver. Since AM or pulse systems have been the most successful and have generated the most interest, the rest of the lidar section of this report will focus on these two types of systems.

Regardless of the waveform employed, an automotive lidar will probably have most if not all of the following features:

1. Semiconductor laser source operating at 850-900 nm,
2. Silicon avalanche photodiode (Si-APD) as the optical power detector,
3. Separate transmit and receive apertures,
4. Inexpensive glass bandpass filter for solar background suppression,
5. Plastic optics,
6. Moderate bandwidth (< 50 MHz) analog receiver section,
7. Digital sampling of return waveform,
8. Scanning or multiple beam system to provide coverage required by the specific mission.

In order to assess how well a typical AM or short pulse lidar would perform, we have constructed lidar range equation models of both based on off-the-shelf laser and detector components and making reasonable assumptions about system parameters. These example systems show a high degree of compliance with the requirements specified in Section 2.4. Their performance is by no means optimized but they can be considered as representative approaches. Table 3.3.1 below summarizes the relevant performance characteristics of the two candidate systems.

	AM System	Short Pulse System
Receive Aperture Diameter	4cm	
System l-Way Transmittance	.6	
Solar Filter Bandpass	250 nm	
Laser Output Beam Divergence	10 mrad	
Laser Wavelength	850 nm	
Laser Power	36 mW peak	2 W peak
Waveform	Sinusoidal AM at 6.25 MHz	Pulse at FWHM=50 ns
Analog Receiver Bandwidth (RBW)	Narrow, centered at f_{AM}	Wide bandwidth = l.5/FWHM
Effective noise bandwidth	FFT RBW- 10 kHz	Same as analog receiver
snr at 80 ft	15 dB	15 dB
Range update rate	RBW/2- 5kHz	PRF=3kHz
Digital Sampling Rate	Nyquist limit at 12.5 MHz	Nyquist limit at 60 MHz

Table 3.3.1 Systems comparison.

The laser powers specified are input powers required per beam and are readily available from commercial lasers. The snr in dB of the received signal has been calculated as a function of range and is plotted in Figure 3.3.1. The laser input power is specified so that a minimum of 15 dB snr is achieved at a maximum range of 25 meters. This snr is required to achieve the .99 probability of detection while keeping the probability of false alarm below 10^{-6}.

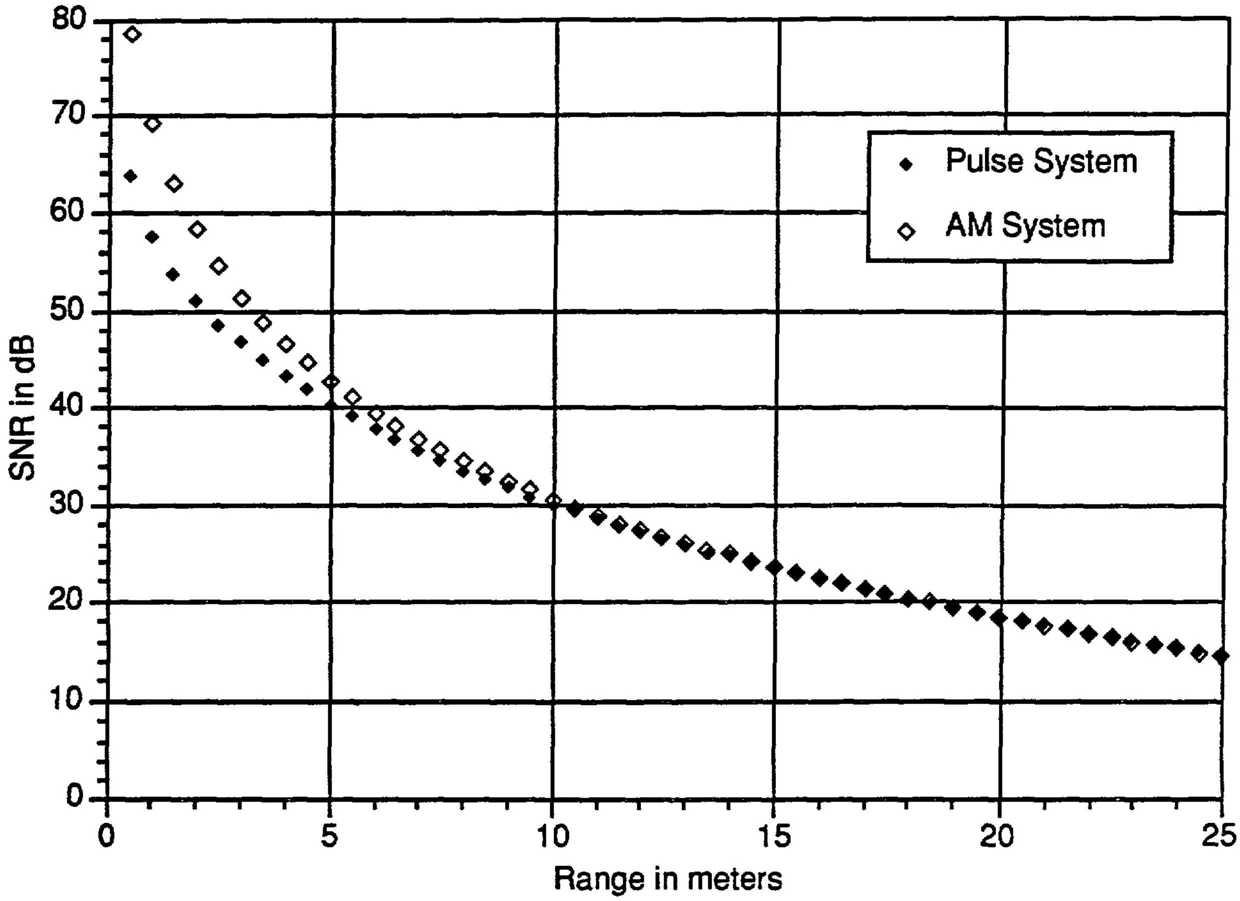

Figure 3.3.1 Signal to noise ratio versus range for both the pulsed and AM lidar systems.

Table 3.3.2 summarizes the expected compliance of the example systems to the preliminary performance specifications of Section 2.4.

Specification	AM System	Pulse Systems
Function	Ranging and derived velocity	
Coverage	1 Sensor per side, 80 ft fore and aft, scanned or multiple beams	
Target Size	Minimum 1 ft diameter at range of 80 A	
Target Velocity	Any allowable (O-65 mph)	
Target Acceleration	$\pm$ 1 g	
Multiple Target Capability	Single target/zone or beam	Multiple targets, range resolved
Platform Velocity	Independent of platform	
Platform Acceleration	Independent of platform	
Relative Velocity	Any allowable ($\pm$ 60 mph)	
Measurement Latency	0.2 ms	0.33 ms
Range Precision	<2ft	< 2 ft
Velocity Precision	TBD	TBD
Probability of Detection	0.99	0.99
Intrinsic Prob. of False Alarm	$< 10^{-7}$	$< 10^{-7}$
Probability of Nuisance Alarm	TBD	TBD

Table 3.3.2 Preliminary performance specifications compliance.

Meeting the requirements for coverage is difficult because the needed range of 80 feet (or possibly longer) requires a low divergence beam for a low laser power system. We believe the performance requirements would be best met by a multiple beam system or single beam system scanned over the positions shown in Figure 3.3.2. This is driven by both the need the keep the laser output power as low as possible and also to maintain a narrow field of view in the receiver, since the largest noise contributor is solar background during daytime.

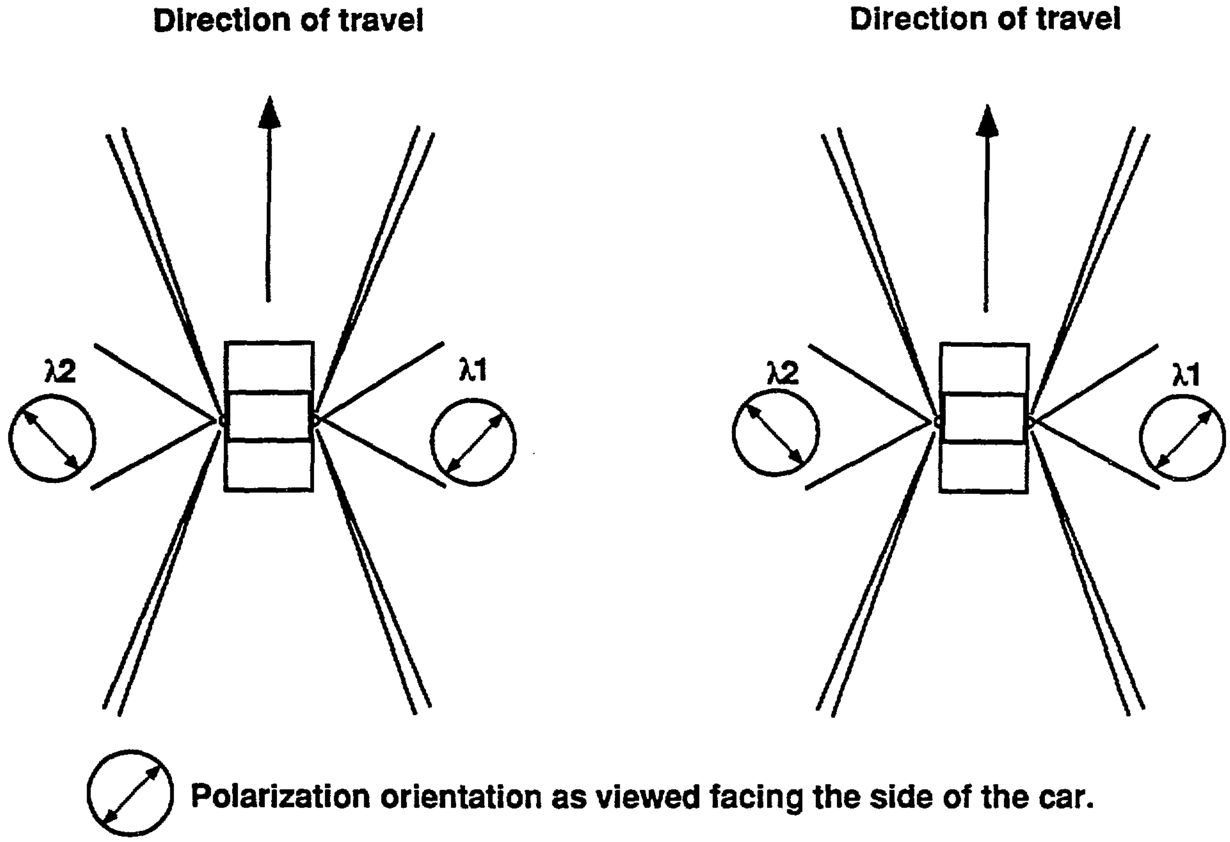

Figure 3.3.2 Conceptual beam coverage of a lidar-based CAS
also showing two interference-reducing approaches.

The measurement latency assumes that the latency is set by the range update rate of the sensor and not by processing speed. Because of the need to reduce interference from other systems, the update rate could be considerably slower than the number quoted above and still meet the latency requirement.

In this case we've chosen to leave the velocity precision unspecified. The range precision estimates assume that the return waveform will be digitally sampled and analyzed. The range precision estimate for the AM system is based on a minimum phase error of 1%, which could probably be exceeded using FFT waveform analysis, although this processing was not specifically modeled for this report. The range precision estimate for the pulse system is based on a detailed digital algorithm developed for a space-based altimeter which takes into account sampling rate, laser temporal pulse width and receiver snr. The probability of nuisance alarm is considered to be extremely algorithm specific and can only be determined through detailed testing and analysis.

3.32 Technologies

All of the technologies required to build a lidar sensor as described here are mature and often available off the shelf. The most important items, the lasers and detector are available as catalog items from EG&G. Tables 3.3.3 to 3.3.5 summarize the laser and detector specifications used to model the example systems. Both operate at 850 nm.

Laser Type	GaAlAs Double Heterojunction
Manufacturer	EG&G Canada
Model Number	C86090E
Maximum CW Power	0.1 w
Spectral line width	4nm

Table 3.3.3 AM laser parameters.

Laser Type	GaAlAs Quantum Well
Manufacturer	EG&G Canada
Model Number	C86083E
Peak Power	10W
Pulse width	50 ns
Response time	< 1 ns
PRF	3000 Hz

Table 3.3.4 Pulsed laser parameters.

Detector Type	Silicon Avalanche Photodiode
Manufacturer	EG&G Canada
Model Number	C30954E
Active Element Diameter	0.8 mm
Responsivity	36A/W
Response time	2 ns
Spectral Noise Density	50 $nA/\sqrt{Hz}$

Table 3.3.5 Si APD detector parameters.

Because the lasers have considerably greater output power than required for the system, we believe it would be quite feasible to split the laser output into three beams and thus avoid the complexity and cost of a scanning system. This does necessitate the use of three separate detectors since the fields of view would not overlap. However the cost of the APDs would probably be less than the cost of a scanning system. However, this has not been examined in detail and could still be traded.

Other important technology elements involve the optics of the system. In order to suppress background sunlight, an optical bandpass filter is necessary to decrease the background induced noise. A multi-layer interference filter would do the best job, but would probably be too expensive. To model this system we chose instead a solid Schott filter glass (RG-9) which provides adequate filtering as long as the detector is not directly illuminated by the sun. We expect that most of the focusing optics would be fabricated (actually molded) from polycarbonate. Polycarbonate is the material used for compact disks and has excellent optical transmittance in the near IR. Also, polycarbonate is extensively used for molded lenses in disposable cameras.

The receiver electronics used for the lidar systems are also very mature. The required bandwidths (30 MHz or less) are moderate by today's standards, and digital sampling at those frequencies is easily done. Less mature is the front-end digital signal processing which we anticipate would be done in either the AM or pulse systems. There is probably some development work to be done in order to optimize the noise immunity of the digital processing. Least mature is the decision algorithm processing to be performed for threat evaluation based on the sensor range and velocity data. We believe that extensive experience with a working sensor would be necessary to mature this technology.

The next five to ten years should see continuing increases in output powers of commercially available semiconductor lasers as well as price reductions in currently available lasers and detectors. We may also expect continuing reductions in the cost of digital signal processing components and ever increasing speeds. These changes will continue to enhance the attractiveness of the Iidar approach to automotive collision avoidance.

3.3.3 Safety Issues

At the power levels anticipated for each of the beams of the system, the only real issue is eye-safety since the average power and pulse energies are well below the maximum permissible exposure (MPE) for skin (14). Here we have to consider a worst case scenario in which someone unknowingly looks directly into one of the transmit apertures of the system. Since the single shot pulse energy is well below the ocular MPE we only need to consider the longer term exposure limit for both systems. For exposure times from 10^{-5} to 10^{+3} seconds the MPE is given by (22)

$$\mathrm{MPE} = 1.8\, C_A\, t^{3/4}\, 10^{-3}\ \mathrm{J\text{-}cm^{-2}} \text{ where } C_A = 10^{2.0(\lambda-.7)} \text{ for } \lambda \text{ in } \mu\mathrm{m}.$$

For a wavelength of about .85 µm, $C_A = 2.0$. We can use the equation for MPE to solve for an estimate of the maximum permissible exposure time (MPET) for the AM and pulse systems. Here we assume that the output aperture area is about 1 cm^2.

The AM system output irradiance is 28.8 mW/cm^2 while the pulse system average output irradiance is 0.24 mW/cm^2 at the maximum PRF of 3000 Hz. The maximum permissible exposure time is given by

$$\mathrm{MPET} = \left(\frac{3.6 \cdot 10^{-3}}{I}\right)^4 \text{ where I : output irradiance in W/cm}^2$$

For the AM system the MPET is 0.2 ms whereas for the pulse system it's 50,000 sec.

Clearly, the AM system cannot be made eyesafe in this worst case scenario without additional safety features or more sophisticated signal processing that would reduce the CW power required of the AM system. In practice, the MPET of the pulse system could be even longer since the prf is unlikely to be as high as 3000 Hz.

On the basis of the peak power required to meet the system requirements and the MPE as specified by ANSI standards, we conclude that the pulse system will be completely eyesafe, whereas the AM system will not be.

3.3.4 Interference Issues

In a widely deployed system the probability of detecting an extraneous source becomes equivalent to the false alarm rate. Therefore we believe a system compliant with the specifications would have to have an intrinsic false alarm rate well below the specified maximum, since the practical false alarm rate will be driven by interference from other vehicles' systems. To distinguish detection of other vehicles returns from the intrinsic false alarm rate, we will refer to the probability of false detection or pfd.

For lidars, we can distinguish two different interference modes, direct and indirect illumination. Since direct illumination will temporarily blind the receiver in addition to producing a false return, additional steps need to be taken to mitigate this effect in addition to whatever signal processing techniques might be applied.

In the case of direct illumination, we need to reduce the optical power on the detector below the saturation level. We can take advantage of geometry limitations inherent in the side-collision avoidance system as shown in Figure 3.3.2. If the systems on opposite sides of the car have sufficiently different wavelengths as to be excluded by the optical filters, then direct illumination from vehicles traveling in the same direction can be very effectively blocked. A strong return would still be registered, but detector saturation would be reduced. Another method which takes advantage of geometry is the use of crossed linear polarizers on opposite sides of the car. This affords attenuation of light coming from vehicles traveling in the same or opposite directions. Either scheme would of course require a uniform standard for all vehicles.

None of these optical techniques could be expected to eliminate false detections due to direct illumination or false returns from indirect illumination, especially from vehicles in the same lane either ahead or in back of the car in question. For pulse systems, the options are limited, since direct detection provides little means of encoding the optical signal. The principal means of reducing the pfd is simply to reduce the duty cycle of the laser while range gating the receiver. Assuming that all systems would operate at the same duty cycle, and that the phases of the cycle would be a uniformly distributed random variable, then the probability of false detection becomes simply

$$\text{PFD} = \frac{\text{Range gate period}}{\text{Duty cycle period}}$$ for a given beam or position of the sensor IFOV (instantaneous field of view).

The maximum range of SO feet gives a maximum time of flight of 160 ns. The range gate needs to be long enough to accommodate the pulse width in the case of the maximum time of flight, so this is about 200 ns. For a pfd of 10" this leads to a duty cycle per beam of 5 Hz,, which is still compliant with the update rate specification. It is possible that a system could experience multiple source interference which would increase the pfd by the factor of the source number. However, the pfd would still be of the same order of magnitude. At 5 Hz a pfd of 10^{-6} results in an average false alarm rate of once every 5.5 hours of operation per beam. Thus we believe a pulse type lidar may be able to meet the pfa/pfd specification under conditions of broad deployment.

AM systems probably have a less straightforward capability to reduce interference. In this case pfd reduction occurs through operation on different modulation frequencies. For the example AM system, the available frequency range is approximately 3.75 MHz based on a minimum modulation frequency of 2.5 MHz. With an FFT-limited resolution bandwidth of 10 kHz **this** leaves only 375 different channels for operation, and the pfd = l/375. It may be possible to reduce this by a combination of frequency agility and averaging of multiple range measurements. Then the pfd would be approximately given by $(l/375)^n$ where n is the number of measurements to be averaged. However, the details of such a signal processing scheme have not been fully developed.

4.0 Processing Technologies

The processor integrated with the sensor forms the heart of the CAS. The processor is responsible for all control of the CAS. It receives all inputs from the vehicle (e.g. speed and turn light indicator state) and setting adjustments from the driver (e.g. sensitivity control, brightness and/or volume control). It interfaces with the display by transmitting the warnings or status of the CAS. It controls the sensor subsystem using parameters which may include timing, mode control, gain control, and data preprocessing. Most of the processor's effort, however, is devoted to the intensive computations required to analyze the raw data received by the CAS sensor in order to interpret it.

These tasks can all be accomplished by a DSP (digital signal processor) which is a specialized processing unit that integrates a number of execution units, memories, ports, controllers, and buses working in parallel.

Shown in Figure 4.0.1 is a typical hardware configuration for the acquisition and processing of sensor data.

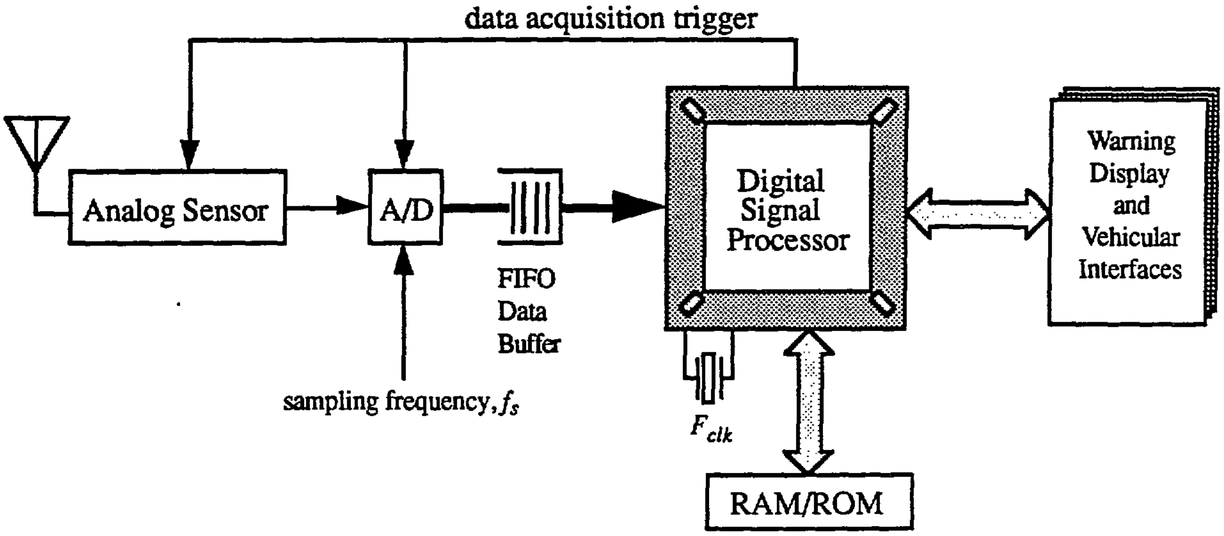

Figure 4.0.1 Block diagram of typical sensor data acquisition and processing hardware configuration.

On a periodic basis, controlled by the clock frequency F_{clk}, the DSP triggers the analog sensor to produce a signal containing information about the scene. This analog signal is then digitized by the analog-to-digital (A/D) converter and buffered with a first-in first-out (FIFO) buffer to create a snapshot of the scene. When the DSP is finished operating on the previous snapshot it is able to rapidly access the data already collected in the buffer. Thus, data is collected while processing of the previous frame is performed in order to eliminate wait states. We will refer to a complete inspection cycle of a snapshot as a dwell. As we will elaborate below, several dwell are often evaluated collectively in order to minimize the effects of measurement noise and ground clutter.

4.1 Processing Requirements

Since the greater part of the DSP's processing is spent interpreting raw sensor data, it makes sense to enumerate some of these associated activities and use this information as a guide to determine the processing requirements. Throughout this evaluation, we will discuss typical processing for a generic ranging sensor.

In general, DSPs are optimized for fast execution of the multiply-accumulate (MAC) instruction. While the MAC instruction constitutes the fundamental mathematical operation of digital signal processing, there exists a wide variety of important tasks the DSP must perform for our CAS application. Nonetheless, the MAC instruction is pervasive throughout the algorithms essential to CAS operation. These include:

Spectrum Estimation

The FFT (Fast Fourier Transform) algorithm is one of the fundamental processing steps used to process raw sensor data in the majority of ranging sensor architectures. For example, in FMCW radar the sensor produces a signal composed of a sum of sinusoids whose individual frequencies are proportional to their respective target's distance. The FFT is able to take this signal and isolate the various frequency constituents to produce a range profile composed of a collection of so-called range gates, or bins.

It has been estimated that 128 range gates (correspondingly, N=256 for FFT, where N is the number of data points to be processed) must be evaluated during each cycle of sensor data acquisition. This will essentially categorize all targets in the effective range of the sensor into one of 128 equally spaced distance intervals away from the sensor. The FFT is computed with O(NlogN) operations, where N determines the number of range gates to be evaluated. While other techniques for spectral estimation such as the ***maximum-entropy*** method and the MUSIC algorithm exist, the FFT is recommended for its processing efficiency. In general, several dwells can be processed consecutively to produce an accurate spectrum estimate in the time it takes to compute these other algorithms which are known for their improved resolution.

Metrics for Preliminary Detection

Once the return energy from the scene is categorized into a set of discrete ranges by virtue of the FFT, the DSP must determine which ranges contain objects of potential interest. This represents the first step in data reduction where a decision is made as to a subset of range gates of relevance and is often a relatively crude estimate. An example of this type of processing is the cfar (constant false alarm rate) detection algorithm which was discussed previously in Section 3.1.3. In essence, the cfar algorithm computes an average of nominal conditions in the scene and declares a detection when one or more binned returns from the scene deviate from the norm by some predetermined threshold.

Tracking and False Alarm Suppression

Preliminary detection, as described in the previous section, is usually a noisy estimate. Often several dwells must be integrated as in an A4 out of N detection scheme. In addition to eliminating false targets, we wish to eliminate targets that are not of interest. These include such things as overhead roadway signs and other irrelevant objects which may appear in the sensor's field of view. To accomplish this, target dynamics are used to characterize and filter out undesirable detections. Algorithms used for this purpose include the Hough Transform and the Kalman Filter.

Hough Transforms

In the computation of range-rate for a particular target, it is inherent that a curve is fit to points corresponding to its time-range coordinates over a local temporal neighborhood of measurements. Using an appropriate parameterization of these curves (nominally as a straight line), a Hough transform (23-26) can be used to perform the fit in a way which ignores outliers. The strength of Hough transforms is a product of its parameter transformation and voting mechanism combination.

Kalman Filtering

Kalman filtering has been shown to be effective in reducing noise with the aid of a state model. A detailed discussion of filters in general can be found in reference (27). The intuition behind a Kalman filter is that the current state estimate is a combination of current measurements and the previous state estimate, transformed to account for system dynamics. The weight given to current measurements is a function of their believed accuracy and how accurately they can be transformed into current coordinates.

Housekeeping Processes

In addition to the core processing algorithms listed above, there will be several other algorithms required which we will refer to as housekeeping. While little detail is given, it should be noted that many other tasks will be performed by the DSP depending on the actual sensor used and the trade-off between digital and analog processing. Examples include adaptive gain control, beam switching for multiple beams, and temperature compensation to name a few. Many of these tasks can be processor intensive and should be accounted for in estimating processing requirements.

4.1.1 Performance Metrics: MIPS, MOPS, MFLOPS

MIPS = Millions of instructions per second; MOPS = Millions of operations per second; MFLOPS = Millions of floating point operations per second. Given these definitions, it is important to note that some instructions are composed of several operations. Since some DSPs have the ability to perform several pipelined operations in parallel, the vendor will often use this to report idealistic values for the above listed metrics. However, in practice, this parallel mode of operation is not sustained continuously. Other practical aspects of DSP operation that impede upon realizing full performance at vendor supplied ratings include: *instruction branching,* which breaks the instruction pipeline (unavoidable); and *slow external RAM/ROM,* which requires the use of wait states (can and should be avoided with faster memory devices).

4.1.2 Fixed Point vs. Floating Point

In the digital domain, three sources of errors are introduced: filter coefficient quantization errors, overflow errors, and round-off errors. Note, however, that each additional bit of wordlength decreases round-off noise power by a factor of four and that coefficient quantization only becomes a serious problem when recursive filtering is performed (e.g. IIR (infinite impulse response) filters are very sensitive to coefficient quantization).

While it has been determined that a l6-bit fixed point processor is adequate for production units, a floating point processor is useful for development and prototyping purposes. In particular, the use of a floating point processor alleviates the designer from many of the tasks associated with making optimum use of the available dynamic range to minimize round-off noise. However, a mere cursory look at the DSP market reveals that floating point processors, in general, cost more than their fixed point counterparts. As with any commercial product cost is a major factor in determining the success of CAS systems in the marketplace. Therefore, it is expected that fixed point processors will be used to lower the unit cost of production units.

A practical measure of a DSP’s merit for the purposes of CAS is the cost to MIPS (millions of instructions per second) ratio which we will refer to as the *alphafactor.*

Most fixed point DSPs have the ability to emulate floating point operations, albeit at the expense of several substitute operations. Thus, a fixed point processor's *alpha factor* would have to be derated to compare to a floating point processor. The derating would, of course, depend on the relative use of floating vs. fixed point processing.

4.1.3 Example Processing Budget Analysis

A need of at least 20 MIPS for algorithm implementation has been estimated on the basis of previous work in this area. Below in Table 4.1 .l, we provide an example processing budget analysis to show an estimate of refresh rate on a presumed processing capability of 20 MIPS.

Processing Task	Number of Instructions	Time Required at 20 MIPS (ms)
Spectrum Estimation	65000	3.25
Target Acquisition	400000	20.0
Target Tracking	200000	10.0
Automatic Gain Control	2000	0.1
Output Data and Timing	200000	10.0
Total Processing Time (ms)		43.35

Table 4.1 .1 Example processing budget for a 20 MIPS DSP.

The processing budget shows that for a 20 MIPS processor, the processing cycle will take 43.35 ms corresponding to a refresh rate of roughly 23 Hz.

4.2 Available Systems (State of the Art)

In order to evaluate the future of DSP hardware, we begin by looking at the state of the art (28). Usability of DSPs is increasing. Manufacturers are producing DSP architectures that are more amenable to C language programming and include more microcontroller functions. Increasingly popular is the ability for current DSP architectures to handle multiple operations in parallel. This enables the instruction decoder to take a complex instruction and simultaneously delegate the associated suboperations to multiple functional units. Some of the latest DSP chips have built-in application-specific features which may or may not be useful for crash avoidance systems (i.e. TI's TMS320C8x with its built-in video controller).

Some of the distinguishing features of available DSPs include:

Instruction/data width

The instruction/data width is a widely used characteristic to categorize the DSP chip. For example, the description "16-bit fixed-point DSP" is used to refer to a class of DSP chips having a 16-bit data bus.

Number of data buses

The number of data buses does not include instruction data buses.

External address range

The external address range indicates an upper bound on the amount of external memory that can be accessed by the DSP.

Hardware-stack levels

An on-chip stack (temporary storage area) is often provided to save the value of the program counter and registers during the service of an interrupt or subroutine. Too many nested subroutines or interrupts may run the risk of depleting the limited stack space built into the processor and can cause the program to crash.

Software Stack

A software stack performs the same function as a hardware stack but is located off chip in system memory. DSPs devoid of a hardware stack will usually have a software stack. Some DSPs are equipped with both a hardware and software stack allowing an overloaded hardware stack to automatically ovefflow into system memory. Larger stack space requirements are becoming necessary to accommodate trends in DSP algorithm development. For a variety of reasons, DSP programming is increasingly performed in a high-level language such as C where code is written as a modular set of procedures. Having a modest amount of stack space allows the developer to maintain code modularity without worrying about the additional stack space required to support nested function calls. In return, sophisticated signal processing algorithms can be developed, maintained and documented with greater efficiency as compared to assembly language programming.

In order to provide a look at where the state of the art in DSP stands today, we have provided a brief list (see Table 4.2.1) of available DSPs (28) that may be practical for the CAS application.

Device	Instruction/Data Width	MIPS	Price ($)	Quantity
Motorola 56002	24/24	20	14.90	100000
Motorola 56166	16/16	30	17.95	100000
TI TMS320C52	16/16	20	14.00	10000
Zilog Z89462	16/16	20	11.00	10000
AT&T DSP1604	16/16	20	13.00	10000

Table 4.2.1 Current examples of fixed point DSPs.

4.3 Future Systems

As a manufacturer of cost-effective CAS instruments one is concerned about the future of DSP chips because they contribute a significant portion of the total system cost and play an important role in system performance. In particular, we are interested in the future cost vs. performance curve as opposed to the future price of a particular level of performance. This cost vs. performance trade-off is important because it is widely believed that improvements in system performance can be attained by “processing gain” (a term used to refer to the improvement of system parameters such as probability of detection through greater algorithmic sophistication, often at the expense of additional number crunching). Consequently, as improvements in DSP hardware continue, the complexity of our algorithms will undoubtedly grow accordingly in order to reap processing gain.

The DSP market is being driven by a number of initiatives with varying degrees of relevance to CAS. For example, there is currently a strong demand for small, low-power DSPs with integrated application-specific functions. This demand is generated mainly by the hand held wireless communications market where size and power are critical. However, as DSP manufacturers continue to address this demand, the benefits of reduced size, weight, or power will not significantly impact the effectiveness or marketability of CAS.

5.0 Conclusions

Both radar and lidar-based CAS have been identified as potentially useful in reducing the number of lane change/merge accidents. This has been accomplished by comparing the performance of hypothetical systems which employ state-of-the-art technology to the preliminary performance specifications that were derived in our Task 4.4 effort. In addition, the processing required for the timely calculation of threat presence and warning issuance has been shown to be consistent with available DSPs. The processor analysis is highly algorithm-dependent. We have utilized our understanding of the required computations to size the DSP for processing speed and memory.

All three technologies are capable of performing at the levels required. However, the costs associated with this level of performance are hard to accurately determine. These costs, of course, scale with the quantities of these systems that are required, which in turn, depends on the acceptance of the CAS. In addition, costs can decrease dramatically in these high technology areas with breakthroughs in materials and manufacturing processes. Both the lidar and radar systems utilize solid state approaches to energy production. These have the promise of being highly compact, robust, and producible at low cost. DSPs are constantly becoming faster and cheaper, and the cost could be reduced even more when they are being produced specifically for the automotive collision avoidance market. There is a significant commonality between various CAS and autonomous/intelligent cruise control products at the sensing and processing level in that they all must detect and monitor targets in specified regions. When a variety of these products starts to become available, this commonality will also lead to "economy-of-scale/scope" reductions.

The performance of the sensor and processing technologies studied in this report will play a major role in determining the utility of various CAS. However, it is clear that the acceptance of these automotive products will depend on their perceived performance by the public. Human factors studies (for example, see reference (29)) have begun to evaluate the effects of missed detections, false and nuisance alarms, and information overload on a number of subjects. The next task addressed under this contract will involve the design of a testbed. This testbed will include systems to monitor the head and eye movements of the drivers (an eye tracker system) and to determine the driving environment around the testbed (two scanning lidar systems). This testbed will accumulate valuable data which will aid in the analyses necessary to help evaluate the effectiveness of future CAS products and thus establish performance specifications.

6.0 References

1. Wassim G. Najm, "A Review of IVHS Crash Avoidance Technologies", Workshop on Collision Avoidance Systems, Virginia, 1994.

2. C.D. Eberhard, P.J. Moffa, S.K. Young, and R.W. Allen, "Development of Performance Specifications for Collisions Avoidance Systems for Lane Change, Merging, and Backing; Task 4 Interim Report: Development of Preliminary Performance Specifications", Contract No. F04606-90-D-0001/0052 and DTNH22-93-X-07022, National Highway Traffic Safety Administration, Washington, D.C., 1995.

3. Stephen .K. Young, Carol D. Eberhard, and Philip J. Moffa, "Development of Performance Specifications for Collisions Avoidance Systems for Lane Change, Merging, and Backing; Task 2 Interim Report: Functional Goals Establishment", Contract No. F04606-90-D-0001/0052 and DTNH22-93-X-07022, National Highway Traffic Safety Administration, Washington, D.C., 1995.

4. Louis Tijerina, Scott Johnston, and Ed Parmer, "Side Crash Avoidance System (SCAS) Test Plan: Test Track Evaluation of the Driver Interface", Draft Report, 1996.

5. Samuel Talmadge, Karen E. Yokoyama, and Greg A. Shreve, "Development of Performance Specifications for Collisions Avoidance Systems for Lane Change, Merging, and Backing; Task 3 Interim Report: Test of Existing Hardware Systems", Contract No. F04606-90-D-0001/0052 and DTNH22-93-X-07022, National Highway Traffic Safety Administration, Washington, D.C., 1995.

6. Merrill I. Skolnik, *Introduction to Radar Systems,* McGraw-Hill Book Co., New York, 1962.

7. David K. Barton, *Radar System Analysis,* Prentice-Hall Inc., New Jersey, 1964.

8. Fred E. Nathanson, *Radar Design Principles,* McGraw-Hill Book Co., New York, 1969.

9. Merrill I. Skolnik, ed., *Radar Handbook*, McGraw-Hill Book Co., New York, 1970.

10. Federal Communications Commission, "First Report and Order, and Second Notice of Proposed Rule Making on Use of Radio Frequencies Above 40 GHz for New Radio Applications", FCC 95-499, Washington D.C., December, 1995.

11. K.W. Chang, H. Wang, G. Shreve, J. Harrison, M. Core, A. Paxton, M. Yu, C.H. Chang and G.S. Dow, "Forward Looking Automotive Radar Using a W-band Single-chip Transceiver", IEEE Trans. on Microwave Theory and Tech., 43, July 1995.

12. D.C. Streit, K.W. Kobayashi, A.K. Oki and D.K. Umemoto, "A Monolithic HBT-regulated HEMT LNA by Selective MBE", IEEE Microwave and Guided Wave Letters, 5, April 1995.

13. H. Wang, R. Lai, D.C.W. Lo, D.C. Streit, P.H. Liu, R.M. Dia, M.W. Pospieszalski, and J. Berentz, "A 140 GHz Monolithic Low Noise Amplifier", IEEE Microwave and Guided Wave Letters, 5, May, 1995.

14. Indira Nair, Alberto M. Guzman, and Louis J. Denes, "Assessment of Potential Biological Effects from Electromagnetic Radiation in the Microwave and Infrared Regions", Task Order No. 2, Contract No. DTNH22-93-D-07007, National Traffic Highway Safety Administration, Washington, DC., 1995.

15. SAE J1113 - Electromagnetic Susceptibility Measurements Procedures for Vehicle Components (Except Aircraft).

16. SAE J1211- Recommended Environmental Practices for Electronic Equipment.

17. Raymond M. Measures, **Laser Remote Sensing,** John Wiley and Sons, New York, 1984.

18. A. Najmi, A. Mahrane, G. Vialaret, and D. Esteve, "Infrared Scanning System for Obstacle Detection in the Automotive Field", Sensors and Actuators A, 41-42, 1994.

19. T. Yanagisawa, K. Yamamoto, and Y. Kubota, "Development of a Laser Radar System for Automobiles", SAE International Congress and Exposition, Detroit, 1992.

20. Manabu Sekine, Tetsuo Senoo, Ikuhiro Morita, and Hiroshi Endo, "Design Method for an Automotive Laser Radar System and Future Prospects for Laser Radar", Intelligent Vehicles '92, Detroit, 1992.

21. Lu Zukang, Tu Dawei, Li Peiyong, Hong Zhi, and Wu Bizhen, "Range Imaging Sensor for Auto-Vehicle Guidance Applying an Optical Radar", SPIE Proceedings 1831, 1992.

22. "American National Standards for Safe Use of Lasers", American National Standard Institute, New York, NY, 1993.

23. J. Illingworth and J. Kittler, "A Survey of the Hough Transform", Computer Vision, Graphics, and Image Processing, 44, 1988.

24. B.D. Carlson, E.D. Evans, and S.L. Wilson, "Search Radar Detection and Track with the Hough Transform, Part I: System Concept", IEEE Transactions on Aerospace and Electronics Systems, 30, 1994.

25. B.D. Carlson, E.D. Evans, and S.L. Wilson, “Search Radar Detection and Track with the Hough Transform, Part II: Detection Statistics”, IEEE Transactions on Aerospace and Electronics Systems, 30, 1994.

26. B.D. Carlson, E.D. Evans, and S.L. Wilson, “Search Radar Detection and Track with the Hough Transform, Part III: Detection Performance with Binary Integration”, IEEE Transactions on Aerospace and Electronics Systems, 30, 1994.

27. Brian D.O. Anderson and John B. Moore, **Optimal Filtering,** Prentice-Hall Inc., New Jersey, 1979.

28. Markus Levy and Anne Coyle, “EDN’s 1996 DSP-Chip Directory”, EDN, 41, 1996.

29. Louis Tijerina, Barry H. Kantowitz, Steven M. Kiger, and Thomas H. Rockwell, “Driver Workload Assessment of In-cab High Technology Devices”, The Fourteenth International Technical Conference on Enhanced Safety of Vehicles, Munich, 1994.

Appendix: Detection Theory

Although the following methodology was developed for radar systems [6-9], it may be used more generally. We begin by discussing the relationship between the probability of detection, the probability of false alarm and the signal-to-noise ratio. Assume that the input noise voltage at the signal receiver is Gaussian distributed, with a mean value of zero and a variance (i.e., rms noise voltage) of Ψ_O. Passing this noise through the receiver's narrowband filter, the probability density of the envelope of the noise voltage output is given by the Rayleigh probability density function

$$P_N(V)dV = \left(\frac{V}{\Psi_o}\right)\exp\left(-\frac{V^2}{2\Psi_o}\right)dV \tag{A.1}$$

When this modulation envelope is passed through an envelope detector, a target detection is considered to have occurred whenever the output voltage envelope exceeds a threshold V_T. Thus the probability of a false alarm due to noise is

$$P_{FA} = P(V > V_T) = \int_{V_T}^{\infty} \left(\frac{V}{\Psi_o}\right)\exp\left(-\frac{V^2}{2\Psi_o}\right)dV = \exp\left(-\frac{V_T^2}{2\Psi_o}\right) \tag{A.2}$$

Note that by raising the threshold voltage, the probability of false alarm may be decreased.

Next, consider a sine-wave signal of constant amplitude A (i.e., a steady-state target model) along with the noise at the input to the filter. The output to the envelope detector now has a probability density function given by

$$P_{S+N}(V)dV = \left(\frac{V}{\Psi_o}\right)\exp\left(-\frac{V^2+A^2}{2\Psi_o}\right)I_o\left(\frac{VA}{\Psi_o}\right)dV \tag{A.3}$$

where $I_O(x)$ is the modified Bessel function of zero order. The probability that the envelope V will exceed a predetermined threshold V_T, which is also the probability that the signal will be detected, is

$$P_D = \int_{V_T}^{\infty} \left(\frac{V}{\Psi_o}\right) \exp\left(-\frac{V^2 + A^2}{2\Psi_o}\right) I_o\left(\frac{VA}{\Psi_o}\right) dV \tag{A.4}$$

This probability increases with increasing signal amplitude and decreasing threshold voltage. This is similar to the relation for the probability of false alarm, where pfa also increases with decreasing V_T. These relationships are graphically illustrated in Figure A.1 below.

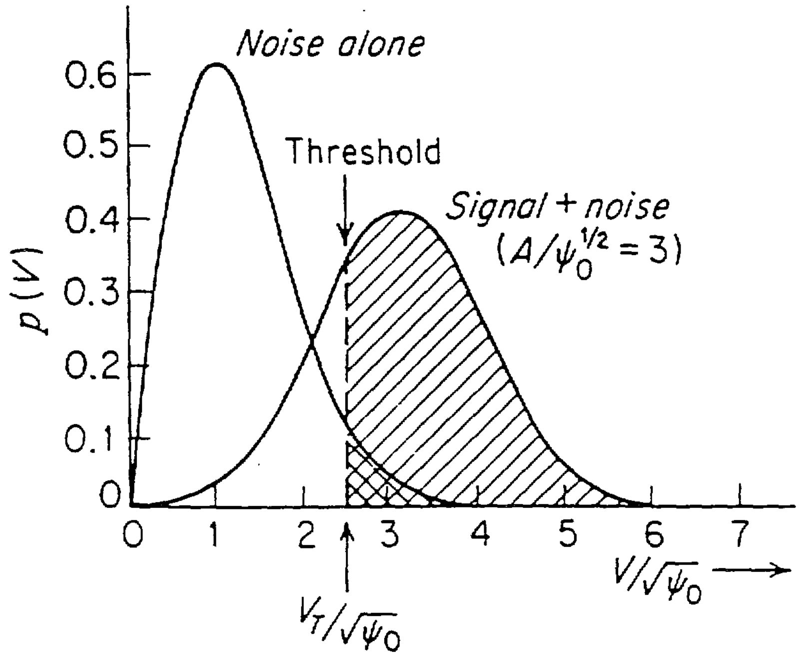

Figure A.1 Probability density functions for noise alone and for signal-plus-noise, illustrating the process of threshold detection.

The probability density for noise alone is plotted with that for signal plus noise. The cross-hatched area to the right of the threshold voltage and under the curve for signal plus noise represents the probability of detection, while the double cross-hatched area under the curve for noise alone represents the probability of false alarm. If the threshold voltage is increased to reduce the probability of false alarm, the probability of detection will also be reduced, whereas if the threshold voltage is lowered to achieve a higher probability of detection, the false alarm probability will also increase.

A final expression for the probability of detection, pd, given by equation A.4, involving just the probability of false alarm and the signal-to-noise ratio may be obtained by noting from equation A.2 that

$$\frac{V_T}{\sqrt{\Psi_o}} = \sqrt{2\ln\left(\frac{1}{P_{FA}}\right)} \qquad \text{(A.5)}$$

and that the signal-to-rms-noise voltage ratio, $A/\sqrt{\Psi_O}$, can be written as a ratio of the signal power to noise power

$$\frac{A}{\sqrt{\Psi_o}} = \sqrt{2\frac{S}{N}} \qquad . \qquad \text{(A.6)}$$

With this, the probability of detection for a given false alarm probability and signal-to-noise can be written as

$$P_D = \int_{\sqrt{2\ln\left(\frac{1}{P_{FA}}\right)}}^{\infty} \xi \exp\left[-\left(\frac{\xi^2}{2}+\frac{S}{N}\right)\right] I_o\left(\sqrt{2\frac{S}{N}}\xi\right) d\xi \qquad \text{(A.7)}$$

where the variable of integration, ξ, is $\xi = V/\sqrt{\Psi_O}$, and $d\xi = dV/\sqrt{\Psi_O}$.

The signal to noise at the receiver due to a vehicle at a given location may be written as

$$\frac{S}{N} = K\frac{P_t G^2 \sigma}{T_o R_{max}^4} \qquad \text{(A.8)}$$

where	S/N	=	signal to noise ratio, snr
	P_t	=	transmitted power
	G	=	transmitter/receiver gain
	σ	=	cross section
	T_o	=	system temperature (used in the definition of noise figure)
	K	=	proportionality constant, specific to the system
	R_{max}	=	maximum sensor range

In general, this expression is appropriate for active systems since the signal-to-noise may be expressed in terms of the transmitted power, a coefficient of proportionality for signal return (a cross section), a detection pattern (the gain pattern), a temperature used in the definition of the noise, and the sensor-to-target distance, R.

Equation A.8 can be utilized to determine the signal to noise ratio for a variety of sensor types. Then the relationships between the signal to noise ratio and the detection probability and false alarm rate can be employed to design a system to achieve the required performance. In this way the required power can be calculated to achieve the necessary detection probability at the desired range with the appropriate beamwidth on the specified target. Details of the system are contained in the system temperature and other system-dependent parameters contained in the factor K.

The considerations presented above do not address the nuisance alarm problem. Unlike false alarms which are due to random noise fluctuations being mistaken for targets, nuisance alarms are true detections based on a threshold-crossing signal. They are termed nuisance alarms because they can cause warnings to be issued when no direct threat exists to the driver. This occurs because the object detected is not in the designed-for detection zone (in space, or velocity, or both) or because it is not a danger because of its size or nature.

There is no comparable theory for understanding nuisance alarms and avoiding them, but some general considerations can be mentioned. Firstly, it is important to confine the emitted energy as much as possible to the desired coverage zone. This may require a scanning or multiple beam system. Also, sidelobes should be reduced as much as possible. Lidar systems with their extremely narrow beam and optical beam formation are much better than radars at accomplishing this.

Another area for consideration is the dynamic range between the returns from appropriate and inappropriate targets. It would be desirable for the energy reflected from valid targets to be significantly higher than from all non-threatening ones but this is not possible. For example, a piece of aluminum foil could produce a very large return for both a lidar and radar system because of its conductivity and reflectivity. On the other hand, a highly aerodynamic or dirty vehicle may produce a significantly reduced signal for a radar or lidar system respectively.

Since, by definition, a nuisance alarm is a detected object, detection approaches alone cannot completely eliminate them. This is why discrimination algorithms have been stressed.