Final Technical Report

Volume I

"Quantitative Characterization of the Vehicle Motion Environment (VME)"

under

Cooperative Agreement No.

DTNH22-92-Y-07319

Conducted for the
National Highway Traffic Safety Administration
U.S. Department of Transportation

Report No. — UMTRI-95-46-1

R. Ervin (UMTRI)
R. K. Gilbert (ERIM)
C. MacAdam (UMTRI)
R. DeSonia (ERIM)
G. Gerber (ERIM)
F. Pont (ERIM)
Q. Holmes (ERIM)
J. McGlynn (ERIM)
K. O'Malley (UMTRI)

December 15, 1995

UMTRI The University of Michigan
Transportation Research Institute

Technical Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.

4. Title and Subtitle
Quantitative Characterization of the Vehicle Motion Environment (VME)

5. Report Date
December 15, 1995

6. Performing Organization Code

8. Performing Organization Report No.
UMTRI-95-46-1

7. Author(s)
R. Ervin, R.K. Gilbert, C. MacAdam, R. DeSonia, et al.

9. Performing Organization Name and Address
The University of Michigan
Transportation Research Institute
2901 Baxter Road, Ann Arbor, Michigan 48109

10. Work Unit No. (TRAIS)

11. Contract or Grant No.
DTNH22-92-Y-07319

12. Sponsoring Agency Name and Address
Office of Crash Avoidance Research
National Highway Traffic Safety Administration
400 Seventh Street, S.W.
Washington D.C. 20590

13. Type of Report and Period Covered
Final
8/1/92 - 12/15/95

14. Sponsoring Agency Code
NHTSA/NRD-50

15. Supplementary Notes

Technical Representative: Mr. Paul Spencer

16. Abstract

The project has addressed the "Quantitative Characterization of the Vehicle Motion Environment (VME)"—seeking to develop a key research tool for building the knowledge base on crash avoidance. By application of this tool, an archival set of data is to be acquired that documents how vehicles are actually being driven in normal usage on U.S. roads. The empirical data in this archive would characterize the trajectories and instantaneous speeds of individual vehicles in the midst of all other nearby vehicles, in everyday traffic.

The work in this project has involved the development of a measurement and processing system for generating and analyzing VME data. A complete ensemble of hardware and software subsystems has been built and subjected to initial trial.

It is clear that the initial technology selected for sensing in this phase of the work—namely, that of laser-based range imaging—is insufficiently mature at present to support the VME program. It appears that the state of the industrial art is well behind that of the raw technological art of laser range-imaging that has been demonstrated in scientific laboratories and for certain military applications. The absence of a commercial market for the peculiar type of laser sensor needed here is rather clearly responsible for the limited industrial capability in this area. The remainder of the VME hardware and software assembly is believed to be entirely utilizeable with substitute imaging technologies. Accordingly, the report examines alternative sensing technologies and establishes that digital CCD technology should be examined for application as the VME sensing medium.

17. Key Words

18. Distribution Statement
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
None	None	110	

Table of Contents

Volume I

Appendices

Volume II

Volume III

1.0 INTRODUCTION

This document constitutes the Final Report on Cooperative Agreement No. DTNH22-92-Y-07319 between the University of Michigan Transportation Research Institute (UMTRI) and the National Highway Traffic Safety Administration (NHTSA) of the U.S. Department of Transportation. UMTRI has engaged this work in collaboration with its subcontractor, the Environmental Research Institute of Michigan, (ERIM). The project has addressed the "Quantitative Characterization of the Vehicle Motion Environment (VME)"—seeking to develop a key research tool for building the knowledge base on crash avoidance. By application of this tool, an archival set of data is to be acquired that documents how vehicles are actually being driven in normal usage on U.S. roads. The empirical data in this archive would characterize the trajectories and instantaneous speeds of individual vehicles in the midst of all other nearby vehicles, in everyday traffic. It is this micro-traffic context, itself, which constitutes the "Vehicle Motion Environment" that must become analyzable in quantitative terms.

The work in this project has involved the development of a measurement and processing system for generating and analyzing VME data. A complete ensemble of hardware and software subsystems has been built and subjected to initial trials. The report is comprised of a presentation of this system, from the viewpoint of its application context (Section 2), the system requirements (Section 3), the design to which hardware and software elements were built (Section 4), and the results derived from testing the system (Section 5).

Based upon the overall experience of implementing and exercising the measurement system, it is clear that the initial technology selected for sensing in this phase of the work—namely, that of laser-based range imaging—is insufficiently mature at present to support the VME program as envisioned by NHTSA. In a macro sense, it appears that the state of the *industrial art* is well behind that of the raw *technological art* of laser range-imaging that has been demonstrated in scientific laboratories and for certain military applications. The absence of a commercial market for the peculiar type of laser sensor needed here is rather clearly responsible for the limited industrial capability in this area. In order to document these judgments regarding the laser sensor, Section 6 of the report presents an overview of the experience obtained with this particular element of the system. Notwithstanding the unsuitability of laser sensing for the VME measurement system, the remainder of the VME hardware and software assembly is believed to be entirely utilizeable with substitute imaging technologies. Accordingly, Section 7 of this report examines the alternative sensing technologies and establishes that digital CCD technology should be examined for application as the VME sensing medium.

The report presents Conclusions and Recommendations in Section 8. In addition, appendices are presented covering the conceptual basis for the design approach, detailed documentation on the measurement and processing subsystems, and the materials developed for gaining eye-safe certification of the laser sensor packages.

2.0 APPLICATION CONTEXT

The central argument for directly measuring the detailed motions of vehicles in normal driving is that the engineering effort to create and evaluate crash avoidance capability on future vehicles will demand it. In particular, the eventual commercialization of so-called "active safety technology" (AST) calls for automotive products that are very well tuned to the dynamic elements of the actual crash-hazard environment through which all vehicles travel. Real drivers tend to anticipate these dynamics from experience and they respond to the exigencies as they develop around them. But the detailed observations we make as drivers are locked away with all of the other so-called "right-brain" skills and adaptations that are vision-dominant and that cannot be meaningfully expressed as a knowledge base for engineering usage. Accordingly, AST advancements and government attempts to evaluate them, will proceed largely in the dark unless this "micro-traffic" context we call the VME becomes usefully measured and quantified.

As conceived here, this need is addressable by means of an instrumentation system that would describe the motions of individual vehicles on a permanent data record. In any of the various applications to be discussed, below, it is clear that a rather faithful recording of each vehicle's trajectory, or motion history, is needed at a sampling rate in the vicinity of 10 Hz. Such sampled motion records are referred to as "track files." At a given road site, one track file would be recorded for each vehicle passing through the scene that is under the immediate observation of the VME system. In addition to trajectory information, the system would also capture the nominal length and width of the vehicle for assessing both the space it takes up and the nominal class of vehicle involved. A fully portable measurement system of this type would be moved from one road site to the next around the country, compiling an archival data set that would eventually represent the near-range behavior of vehicles operating in traffic in the U.S.

Since the measured track files would all have a common time base, later processing of these data can determine the inter-vehicular relationships that prevailed during the measurement. Such "enriched-variable files" can each support new statistical analysis if inquired simply as another layer of data. Alternatively, a set of direct track file data plus an ancillary file of derived variables could be generated to support the simulation of an AST prototype system as an overlay on the "truth environment" comprised of VME data.

An engineering characterization of the VME will require that measurements be made at selected sites for a modest period of time—say a month or so at each site. The full sample of such measurements must cover a representative sample of sites covering geographic, climatic, road design, illumination, driver, and traffic factors. At a given road site, each motion and space variable must be quantified from one instant in time to the next so that, eventually, data are collected providing statistical distributions of these variables representing the vehicle operations within which AST packages would be deployed. Altogether, such an archive would constitute a massive data resource and would require a sustained commitment for its acquisition and maintenance—not unlike the commitment that has attended the compilation of the computerized accident record.

Without VME data, it is felt that the process of refining collision warning and intervention systems will be remarkably empirical in nature and thus quite handicapped as an engineering endeavor. The empiricism will derive from the simple fact that the precrash environment remains utterly unquantified. Thus, the only way one can tell if a given sensor/processor package is any good, under the current state of affairs, is to take it out on the road and try it. But wherever one tries it, the inter-vehicular variables at the time of testing will be unknown and unrepeatable in any controlled sense—thus making it difficult to relate the package's performance to the condition variables. Given that the population of drivers exhibit a substantial level of random variation in all control actions, attempts to simulate this application environment will always lack validation until some robust form of "truth data" is brought forward through a direct-measurement characterization.

The basic problem is that we have essentially no information that is both quantitatively and statistically representative of the headways, lateral clearances, angles of approach, time spacing between vehicles, or the correspondence between these inter-vehicular variables and the steering and braking accelerations which are driver-induced in response to this motion environment. Thus, we are without definitive data on an exceedingly complex application environment toward which a large industry around the world is now targeting a vast array of new technology, promising crash avoidance countermeasures.

The extent of the need for VME data can be seen upon consideration of the challenge in AST system development. The central observation, confirmed now by some industry engineers who have begun to work on active safety packages, is that the detection of full-blown, fast-closing collision threats is not too difficult if the system waits long enough to make a decision. But then, the time-to-respond may be intolerably short. Many sensing technologies, even with crude processing algorithms, can tell a bona-fide crash-in-the-making when it is well-developed and more or less inevitable. The hard part is to create sensor/processor systems that can discern the "probably-harmless" inter-vehicular actions from the "very-likely-harmful" events early in the time sequence. Clearly, since candidates for crash-interaction develop around each motor vehicle hour after hour, day-in and day-out, the opportunities for false alarm are many, indeed. The suitable active safety technology must accomplish the remarkably complex task of accepting the many thousands of episodes which are, indeed, benign while not ending up in such a mathematical stupor that the bona-fide collision threat is missed or its detection is delayed beyond the minimal time window needed for safe intervention. On the expectation that frequent false alarms and, worse yet, false control interventions, will simply make such systems unacceptable, the achievement of high levels of "active safety intelligence" seems a requirement. But the engineering of such intelligence into these products appears, in turn, to require an accurate targeting of the technology to the complex motion environment as it really prevails. Such a task, in turn, requires that "the target" be representatively quantified.

However industry may use such quantitative data for product planning and development, government may be disposed to employ the VME data for such purposes as identifying opportunities for crash avoidance countermeasures, proving AST concepts at a preliminary level, and evaluating specific system designs by subjecting them to VME sequences that have been selected from the archive for use in repeatable and statistically meaningful examinations of product performance. A "standard" evaluation sequence might emerge by which industrial developers of technology can communicate with government regulators, and vice versa, perhaps eventually even using a VME data sequence to develop product standards covering certain "macro" aspects of safety performance.

In this context, a national archive of VME data has been viewed as analogous to the archive of accident data, itself. That is, just as we have used the police-reported accident record to document our national crash experience and, in turn, to help in developing "passive safety technology," so the VME data record would document our national everyday-driving experience in terms that would help develop, evaluate, and perhaps regulate an active safety technology.

3.0 SYSTEM REQUIREMENTS

The goal of this project is to demonstrate that the VME concept is practical. The goal will be met by building a measurement and analysis system in the form of an operating prototype and demonstrating its utility through a trial data collection, processing, and evaluation effort. The objectives are as follows:

1. to design a measurement system that:
    a. meets the performance requirements (specified later)
    b. is portable and deployable at any road site
    c. is as economical as possible, given the performance requirements
    d. is modular, allowing tailored coverage of any desired road section by means of a set of sensors, each imaging a roadway are which overlaps the next.
    e. is safe, in terms of constituting a fixed object at the roadside
    f. yields track files in a manner that maximizes their utility for analysis. (A track file, defined later, documents the motion of an individual vehicle through the sampled road segment.)

2. to fabricate the VME-MS with sufficient modules to support a significant road site demonstration of the method

3. to design the software-based processing system, the VME Data System (VME-DS), which operates on track file data to enable various forms of meaningful analysis

4. to demonstrate the operation of the VME-MS at a public road site, generating track file data for all vehicles passing through the site over approximately 50 hours of actual traffic operations

5. to process the track file data through the VME-DS, demonstrating the utility of the software package for conducting all of the analysis options

The VME System is composed of two subsystems: 1) the VME-Measurement System (VME-MS) and 2) the VME-Data System (VME-DS). The functional and performance requirements for the VME-MS and VME-DS are presented in Section 3.1 and 3.2, respectively.

The VME-MS shall be designed for roadside deployment for the purpose of producing, in real time, track files on all vehicles passing through the system's field of measurement. The track-file data shall be processed, off line, by the VME-DS producing refined estimates of each vehicle's trajectory. The VME-DS processed vehicle location and heading accuracies shall be as specified in Figure 3-1, Vehicle Motion Data.

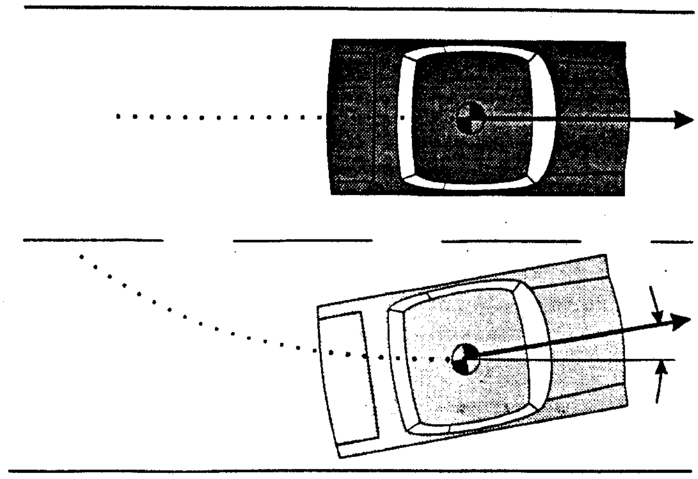

Track File - Each Vehicle

Time sampling rate	10 Hz
Vehicle Rectangular Shadow (LxW)	+/- 3 inches
X-Y coordinates of the vehicle's centroid vs. time	+/- 6 inches
vehicle yaw angle vs. time	+/- 2 degrees

Figure 3-1. Vehicle Motion Data

3.1 VME-MS Requirements

The Vehicle Motion Environment Measurement System (VME-MS) shall be designed to accurately measure vehicle trajectories at roadway sites and produce track files for all the vehicles passing through the system's measurement field-of-regard. The vehicle track files (body centroid position and body yaw angle as a function of time) will be generated in real-time and recorded at field sites. The VME-MS shall be designed in a modular fashion such that the VME-MS measurement field-of-regard can be increased through the simple addition of modular, measurement units. Imaging sensors, operating at 10 frames/sec, and commercially available general purpose computer processors for system control and data processing, shall acquire data on all vehicles within the measurement field-of-regard for a single VME-MS module. Co-located with each sensor shall be a video camera and its associated storage media for providing backup video whenever the system is operating. The VME-MS Performance Requirements are summarized in Table 3-1.

To the maximum extent possible, the VME-MS shall be designed following good engineering practices and built using commercial off-the-shelf components. The VME-MS shall be designed in such a manner that the roadside deployment procedures do not require any sophisticated, special purpose tools.

Table 3-1. VME-MS Performance Requirements

Coverage Area per Sensor Station	60 feet wide x 200 feet long, minimum coverage
Number of Sensor Stations	3 minimum; extendible to 15
Deployment	Along roadway, around intersections, at freeway on/off/merge lanes, etc.
Measurement Sensor	Laser-based imaging sensor mounted on 100 ft high, portable tower
Control and Data Processors	Commercial general-purpose computers (80486/68040 class)
Output Data	Real-time generation of vehicle track files for all vehicles with in each sensor station's field-of-view Video camera recordings simultaneous with 3D-Laser sensor operation; data shall be "tagged" in such a manner that the video and track file data can be easily correlated
Storage Capacity	Sufficient for 24 hour, continuous operation
Storage Media	Removable, commercially available
Temporal Sampling Rate	10 frames/second
Data Accuracy	Vehicle centroid track ±6 inches Yaw angle ± 2°. @ 120 feet slant range Yaw angle ± 8° @ 300 feet slant range
Coordinate System	Data referenced to Roadway Surface in Cartesian coordinates
Operation	Up to six (6) months deployment in the field; 24 hours continuous operation between service for removal of recordings and test of system performance Under all weather conditions down to 1/4 mile visibility and winds up to 75 mph Temperature/Humidity Ranges, prevailing conditions at any location with the continental United States at any time of the year

3.2 VME-DS Requirements

The VME-Data System (VME-DS) receives VME-MS Track Files that must be processed to reduce the data magnitude for archival storage (see Figure 3-2). Data will ultimately be stored in the "Data Base of VME Truth" based upon one of two sampling approaches. On the top branch, a continuous sampling produces the basic data set (i.e., all track files.) On the bottom branch, the system retains raw image data only from the infrequent near-miss and accident events that do occur, allowing detailed scrutiny of these data in their original, highly-defined, image format. To implement the right-hand branch, a simple incident detection algorithm must be employed to capture the anomalous events. A detector algorithm might look for an impending crash condition, for example, by noting when the ratio, DV/C, exceeds a threshold level TH, where DV is the closure or relative velocity between two nearby vehicles and C is the instantaneous clearance distance between the two.

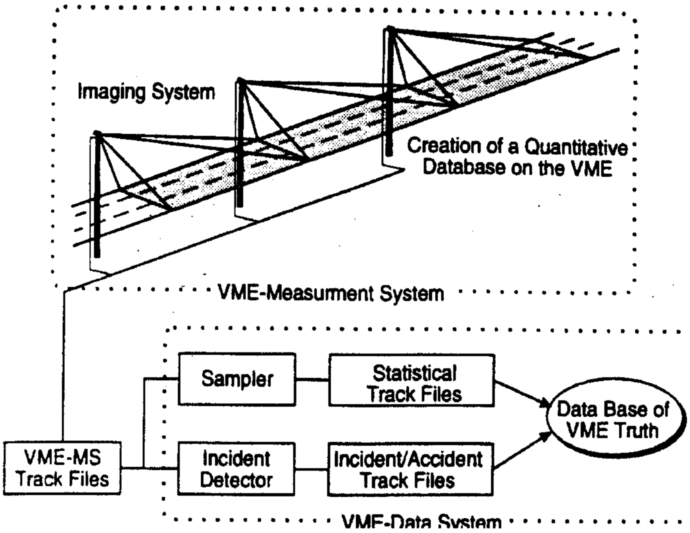

Figure 3-2. Creation of a Quantitative Database

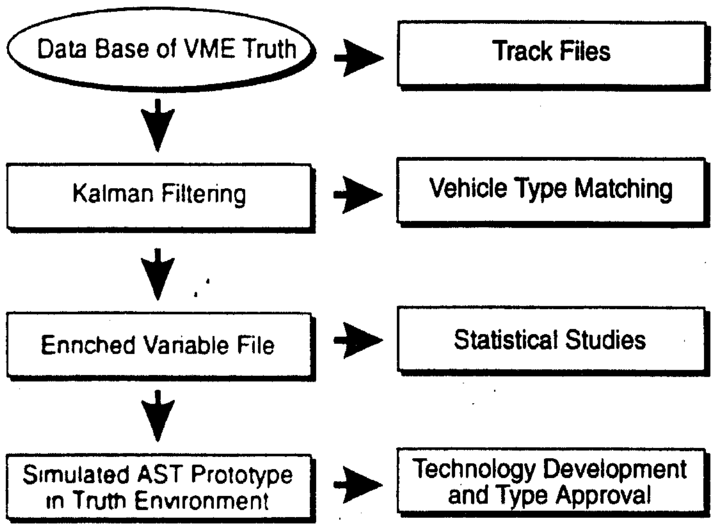

Figure 3-3. Possible Processing of Database

Shown in Figure 3-3, the generated data base can be further processed, down the left column, in order to support products of special interest, down the right column. As is, the data base can be scrutinized at the track file level to derive inter-vehicular variables such as approach velocities, attack angles, instantaneous clearance, etc., as a function of time. Kalman filtering on track file data can produce information for precise type-matching of vehicles and can also derive other variables that were not directly observed such as the driver's steering wheel and braking inputs, body sideslip angle, etc.

The "enriched-variable file" can then yield new statistics if inquired simply as another layer of data, or can directly support analyses in which AST prototype systems are simulated "in the truth environment." For example, a run-off-road warning system being developed in the future may employ a signal representing steering wheel angle as well as sensory signals showing vehicle position and heading angle relative to lane edges. If the VME data file is to fully support the simulated operation of such a system, it must contain these continuous variables in the file.

3.3 VME Interfaces

The VME has two principle, external interfaces, one with the VME-MS deployed at the roadside and the second with users of VME-DS products and services. The interface requirements for the VME-MS are specified below. Section 4.2 discusses the VME-DS interface.

3.3.1 VME-MS Roadside Interface

The road-side deployment of an individual Sensor Station is illustrated in Figure 3-4. The imaging sensors view the roadway from the top of a 100 foot high tower and a weather-proof enclosure at the base contains all electronics needed for control, processing and data recording. The tower is portable for deployment to other sites and is protected by an appropriate vehicle barrier. A security fence is also necessary to prevent persons from gaining unauthorized access to the tower.

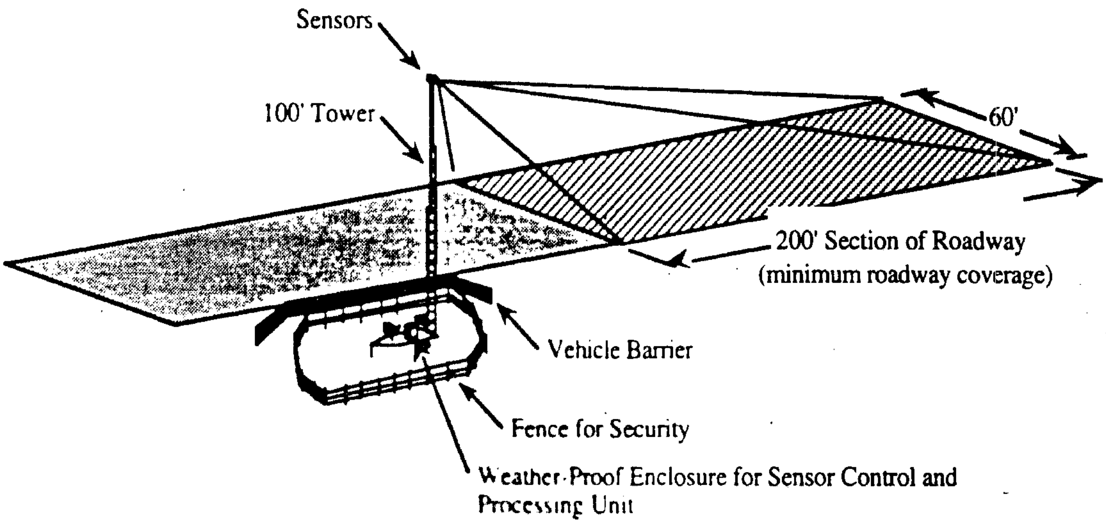

Figure 3-4. A VME Sensor Station is Portable for Deployment at Selected Road-Side Sites.

For each deployment site, the site plan should be submitted to the Federal Aviation Administration (FAA) to ensure compliance with regulations on height of structures. Also, approvals and permits at appropriate levels (e.g., state, county, municipality) must be obtained prior to deployment. Consult local authorities to ascertain the necessary approval chain because the process will vary from location to location. Any height and setback variances must be obtained prior to the deployment of the system.

This approval process can often be time consuming, so plan ahead and begin the approval process during the early planning stages.

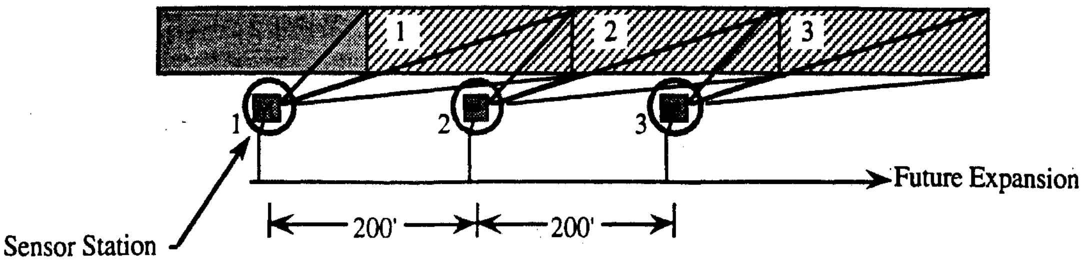

(a) Three Sites Spaced at 200 ft. for Contiguous Coverage of Roadway (600 ft. Long)

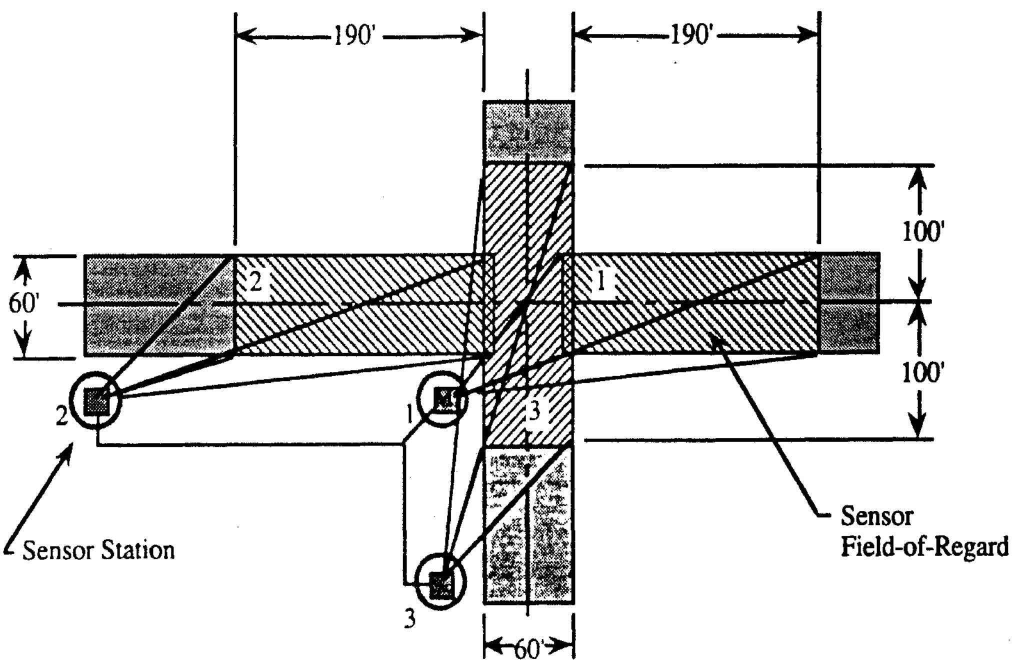

(b) Three Sites Spaced Around Intersection

Figure 3-5. A VME Sensor Station is Portable for Deployment at Selected Road-Side Sites

The VME-MS shall be designed and built to be eye-safe when deployed in a public setting along roadways or at intersections.

Multiple 200-foot sections of freeway or roadway intersections shall be covered by deploying multiple Sensor Stations. Three Sensor Stations will be provided for this first implementation of the VME-MS with representative deployments illustrated in Figure 3-4. A number of "control points" shall be deployed within the measurement field-of-regard. The retro-reflectance characteristics of the control points shall be sufficient to produce a very reliable signal under the worst visibility conditions. The number of and location of the control points shall be selected in such a manner to minimize system complexity and sufficient to satisfy the VME-MS Performance Requirements list in Table 3-1.

4.0 SYSTEM DESIGN AND IMPLEMENTATION

The VME system hardware and software components are described in the following Sections 4.1 and 4.2.

4.1 VME-Measurement System Design and Implementation

The essential elements of the VME-MS physical design and implementation are presented in the following subsections. Specific details such as engineering drawings and an operator's manual are provided in the appendices. This section is divided into the following five (5) subsections.

System Architecture - describes the modular nature of the VME-MS design and the allocation of major functions to hardware and software elements. The modular structure of the VME-MS design allows for the addition of an arbitrary number of Sensor Stations, thus providing the capability to monitor any desired length of roadway.

Sensor System Design and Implementation - describes the major hardware and electrical elements that comprise the LASAR DatacameraTM with particular emphasis on the Laser Sensor Head, which was an unanticipated new design for the vendor. The intent was to leverage off the design of an existing commercial product, but using the increased power and higher scan rate requirements imposed by the VME Program

Mechanical Design and Implementation - describes the hardware elements and their functional decomposition for a Sensor Station, which is the basic building block of the VME-MS. The construction of the VME-MS emphasized the use of existing commercial equipment, or commercial equipment that was slightly modified by the vendor. To minimize the costs associated with fabricating multiple copies of a Sensor Station, very few custom piece parts were designed and fabricated.

Electrical Design and Implementation - describes the implementation of the real-time computing, control and communications elements using conventional, PC-hardware, related peripherals and local-area network equipment. This approach allows for the easy, and inexpensive, upgrading of the VME-MS capabilities as the commercial industry introduces higher performing hardware.

Software Design and Implementation - describes the message-based, real-time computing system developed for the VME-MS. The software is a blend of commercial software packages for standard computer operating system functions and local-area networks and custom software for high-level system control (i.e., at the Sensor Station level) and the image and data processing that is sensor and application specific.

4.1.1 VME-MS System Architecture

By deploying a 3-D imaging sensor on a sufficiently high tower at intervals along a freeway, as illustrated in Figure 4-1, 3-D imaging technology can be exploited to sense the positions and sizes of vehicles which are traversing each sensor's field-of-view. Each 3-D sensor and associated tower, transporting trailer, and supporting electronics comprise a sensor station. The sensor stations are interconnected via the communication network and the data produced by the individual stations are "fused" together in real-time to produce vehicle track files which appear to the data analyst as coming from a single "virtual" 3-D sensor covering the combined viewing area. The vehicle track file data is stored on magnetic tape for subsequent (off-line) processing by the transportation research community. Potentially, the vehicle track file data could be sent via large area network (e.g. Internet) for immediate processing, if there is need for more immediate analysis.

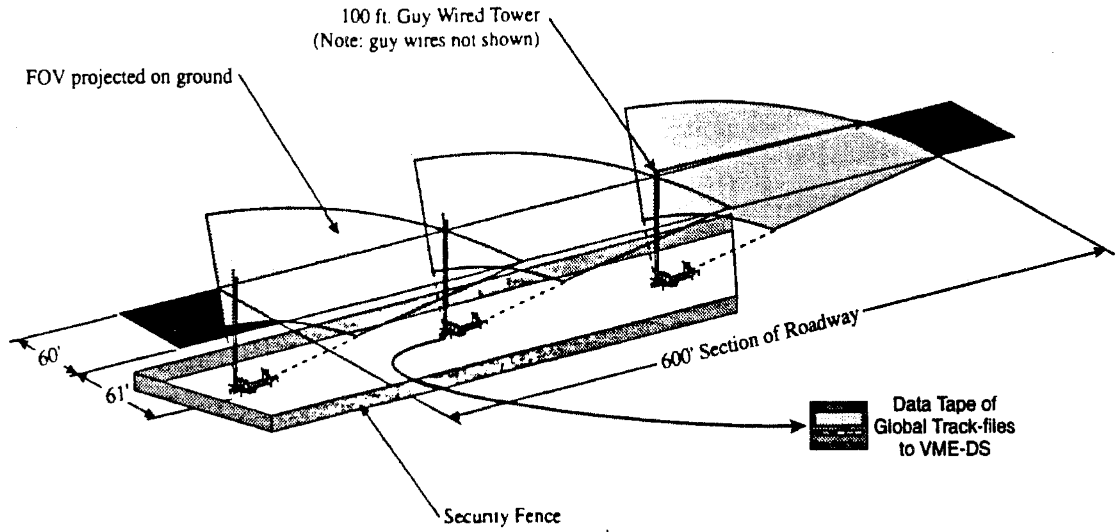

Figure 4-1. The Demonstration Deployment, A Straight Freeway Configuration.

The viewing geometry was chosen to maximize the quality of track file data within the operational constraints imposed by the roadway environment. The active area or field-of-regard that must be covered by a single sensor station is 200 feet long by 60 feet wide. Figure 4-2 depicts the viewing geometry both from a plan view and side view. Pertinent values of angles and distances are indicated on the figure. A 100 foot tower was chosen to minimize mixed pixel effects and vehicle obscuration while aiding laser safety. Transportable (and thus guyed) towers were selected for the demonstration phase of the VME-MS program and must be offset 50 feet from the roadway edge. Other sensor placement options are available for more permanent locations including non-guyed towers, or use of existing structures such as buildings that might be found near intersections.

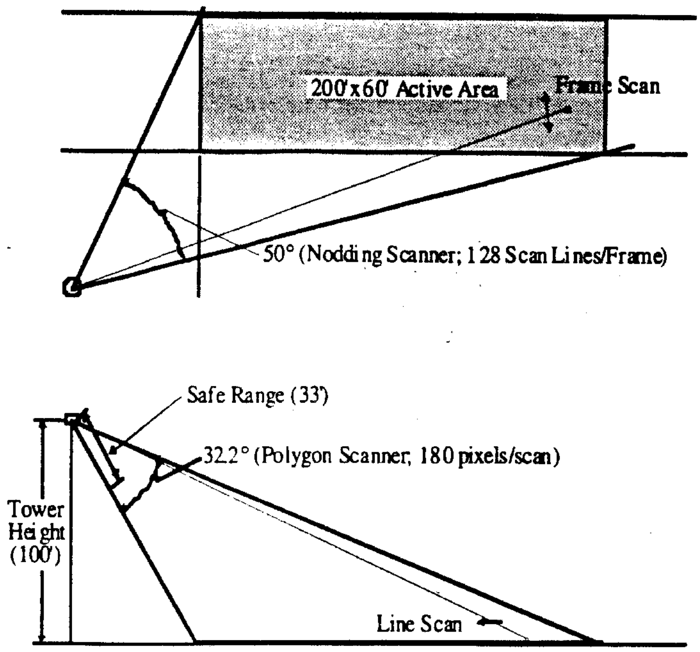

Figure 4-2. Viewing Geometry of the VME-MS

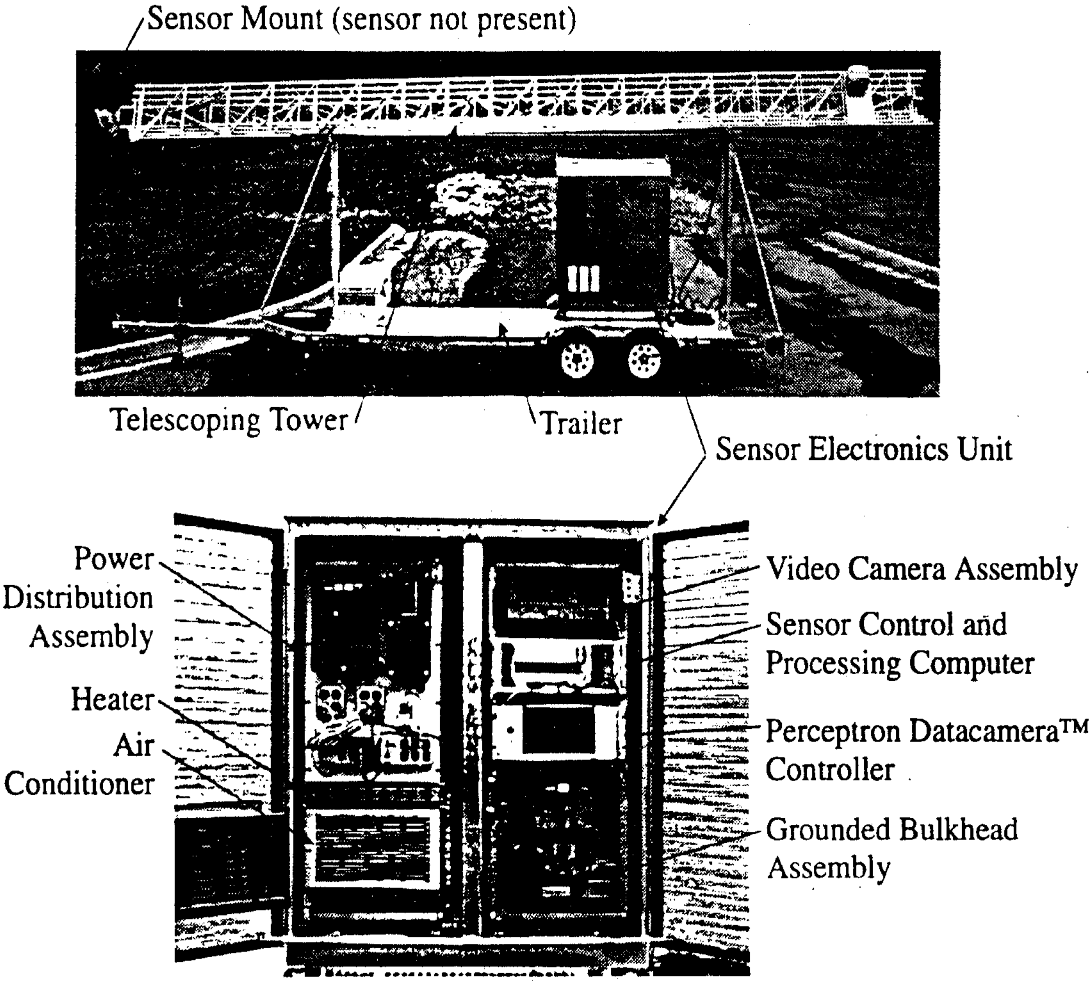

Figure 4-3. Assembled VME-MS Sensor Station

Figure 4-3 shows a fully assembled VME-MS Sensor Station. Each sensor station has associated supporting electronics mounted in an environmentally-protected enclosure (the Sensor Electronics Unit or SEU) on a trailer at the base of the tower. One of the Sensor Stations is augmented with additional electronics and is designated as the Master Electronics Unit (MEU). The MEU is identical to a Sensor Electronics Unit except for the addition of the Master Control & Processing Unit (MCPU), an industrial-grade PC with external storage devices, which shares an environmental enclosure with SEU electronics for packaging economy. A conventional CCD video camera is mounted on the tower (boresighted with the laser radar) and provides video imagery input to a video tape recorder for diagnostic analysis.

The VME-MS modular architecture (diagrammed in Figure 4-4) is comprised of both distributed and centralized functional elements. The computationally-intensive and sensor site-specific tasks are performed at each sensor station.

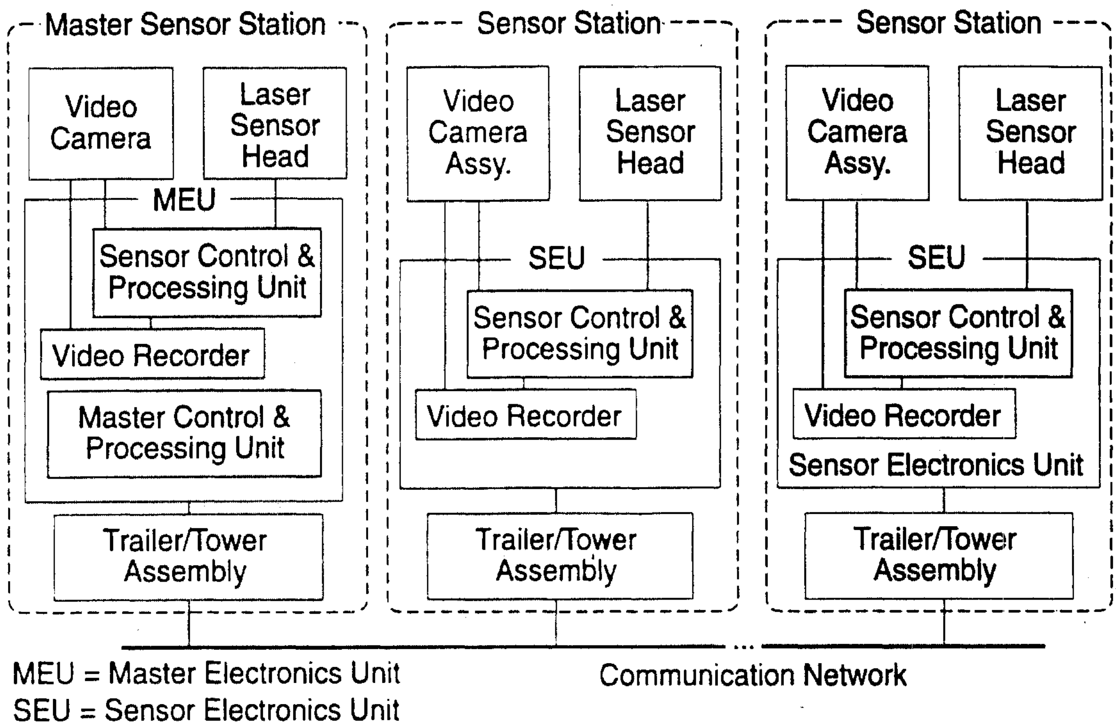

Figure 4-4. VME-MS Architecture

The architecture is readily extendible to increase roadway coverage through the addition of more Sensor Stations to the communication network. The MCPU processing capacity will limit the maximum number of Sensor Stations to about 15. Further increase will require the addition of one or more MCPUs to the system. The design supports the addition of a "super MCPU" which ties together groups of MCPUs for extended highway coverage. The communication network design (a standard 10Base-T Ethernet network) requires the addition of repeaters to increase the linear coverage beyond the currently planned 600 feet. Space has been allocated in the electronics enclosures for a network repeater (when required). The network topology may be extended to cover much larger areas using proven Local Area Network (LAN) techniques. Transfer of the track file data from (possibly distributed) MCPU units via a Wide Area Network (WAN) (e.g. Internet) to a centralized data collection facility is a future growth path supported by this architecture.

Functional Organization within an Electronics Unit

The MEU contains the MCPU, which provides centralized control and coordinates the tracking-data acquisition activities of all sensor stations (including the one where the MCPU is resident) via network packet communication and is unique to the Master Sensor Station. The MCPU provides the following (centralized) functions: 1) processing of the vehicle tracking data from multiple sensor stations into a single data stream (for archiving); 2) data archiving (via a magnetic tape drive); 3) storage of VME-MS system software (via a magnetic disk unit); 4) power-on initialization, setup, and control of the sensor stations; 5) archiving of data from sources external to the VME-MS that may be available (e.g. road-surface condition sensors); and 6) remote status monitoring (via modem interface to a telephone line or cellular phone network).

Every Sensor Station contains a Sensor Control and Processing Unit (SCPU) which provides the following dedicated functions: 1) power-on initialization, setup, and control of the Laser Ranging Camera System (LRCS) and the Video Camera System; 2) processing of the image data from the LRCS into track feature data; 3) transmission of the track feature data to the MCPU; and 4) recording video data.

Communications

Distributed computing elements are linked via a communications network using packet communication between a central control unit (the Master Control and Processing Unit) and the Sensor Control and Processing Units (SCPUs). Figure 4-4 shows the communication network architecture. The packet-based communications exploits commercially available networking products (e.g. Novell's Personal Netware) by utilizing their underlying packet communication capabilities (IPX) developed to support their file servers. Each SCPU also controls distributed elements (the Laser Sensor Controller and the Video Recorder Assembly) via RS-232 communication links.

Control and Processing

The processing unit functions may be broadly divided into Image and Data Processing functions and Real-Time Control functions (including network communications). The VME-MS is designed to operate as a turn-key system once initial site setup is complete, requiring no user involvement other than replacement of the magnetic tape (4 GByte DAT) when full. During the site setup phase, a diagnostic and maintenance PC is attached to the communication network, which provides a user interface to the VME-MS and supplies diagnostic and set-up software tools. The diagnostic computer can monitor the VME-MS system during normal operation or, in a VME-MS standby mode, individually access each of the Sensor Stations and remotely exercise local functions, capture and transfer images for live display, run diagnostics, and alter setup information (such as sensor location in the global coordinate frame).

During normal operation (and following initialization), the VME-MS is "event-driven" by the servicing requirements of the laser radar output data stream. Real-time control is "open-loop" as no output results feed back to alter the input state. Output data from a Sensor Station observing an "upstream" vehicle is passed to the "downstream" sensor station via the MCPU as "heads-up" notice of an approaching vehicle and to better initialize the downstream station's tracking coefficients for that vehicle. Time synchronization of the sensor stations is maintained by the MCPU via inquiry of the time of receipt of a special broadcast packet periodically sent by the MCPU; time corrections to counter individual SCPU clock drift are sent by the MCPU to each SCPU as necessary. The MCPU is "message-driven" by the servicing needs of the Sensor Control & Processing Units to off-load vehicle track data once a vehicle has left the sensor field-of-view or to request a vehicle identification number when a new vehicle enters the field-of-view, for example. Figure 4-5 shows the data flow within the VME-MS system.

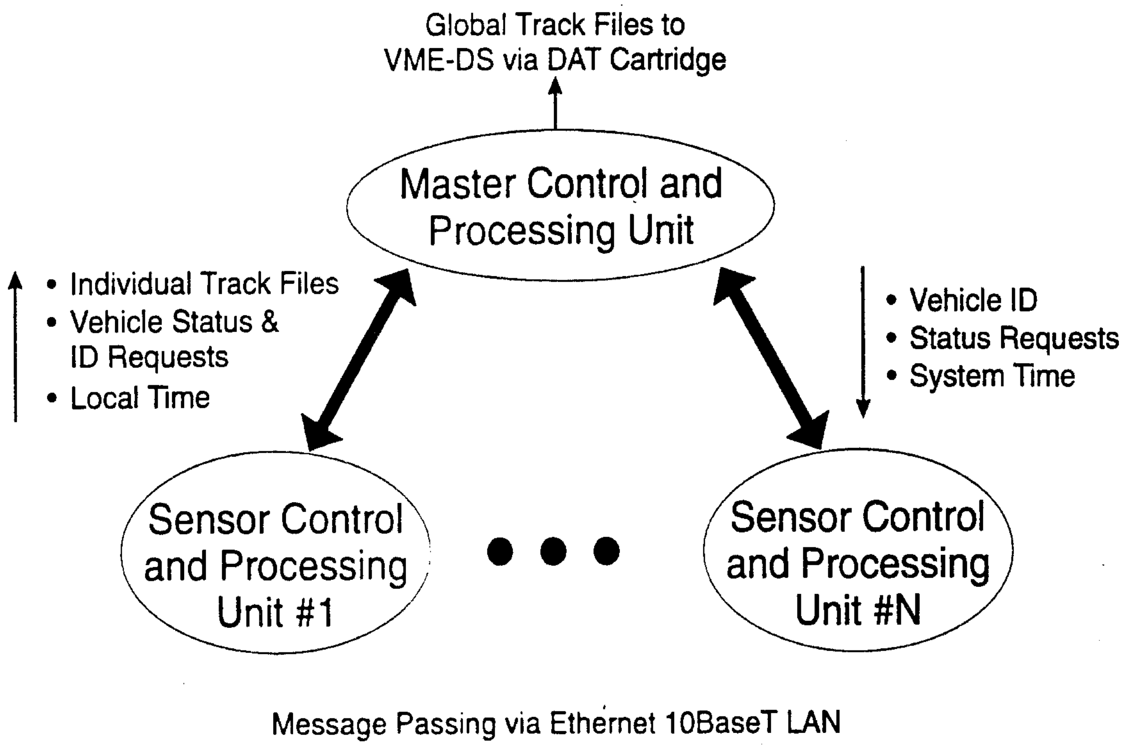

Figure 4-5. SCPU <-> MCPU Communications

Both control and processing functions were implemented using a 32-bit version of PolyForth™ (Forth Inc.) running in protected mode, but making use of IPX (Novell) packet transfer capabilities and ASPI SCSI driver (DAT tape) interface software operating under DOS. Since the requirements for packet communications and tape transfer are relatively infrequent, the time overhead necessary to switch the CPU to virtual 8086 mode is insignificant and the cost-benefits of using inexpensive, commercially-developed software developed for the PC marketplace could be achieved.

Figure 4-6 illustrates, at a high level, the sequence of operations which occurs as 3-D image frames are processed (by multiple sensor stations). After acquiring an image frame, each Sensor Station performs single-frame processing on the range image to detect vehicles and extract vehicle features. This (non-image) data is combined with data collected from previous frames in order to refine the vehicle features and track information is temporarily stored. When a vehicle leaves the sensor field-of-view, the time-tagged vehicle positions, headings and other attributes (e.g. size) are sent to the

Master Sensor Station (MCPU) where the data is temporarily stored until the track data for a particular vehicle has been received from all (three) sensor stations. The track data is then combined into a single track file (referenced to a Roadway Cartesian Coordinate System) and archived on magnetic tape. When full, this tape is transferred to the data analysis center for track file processing.

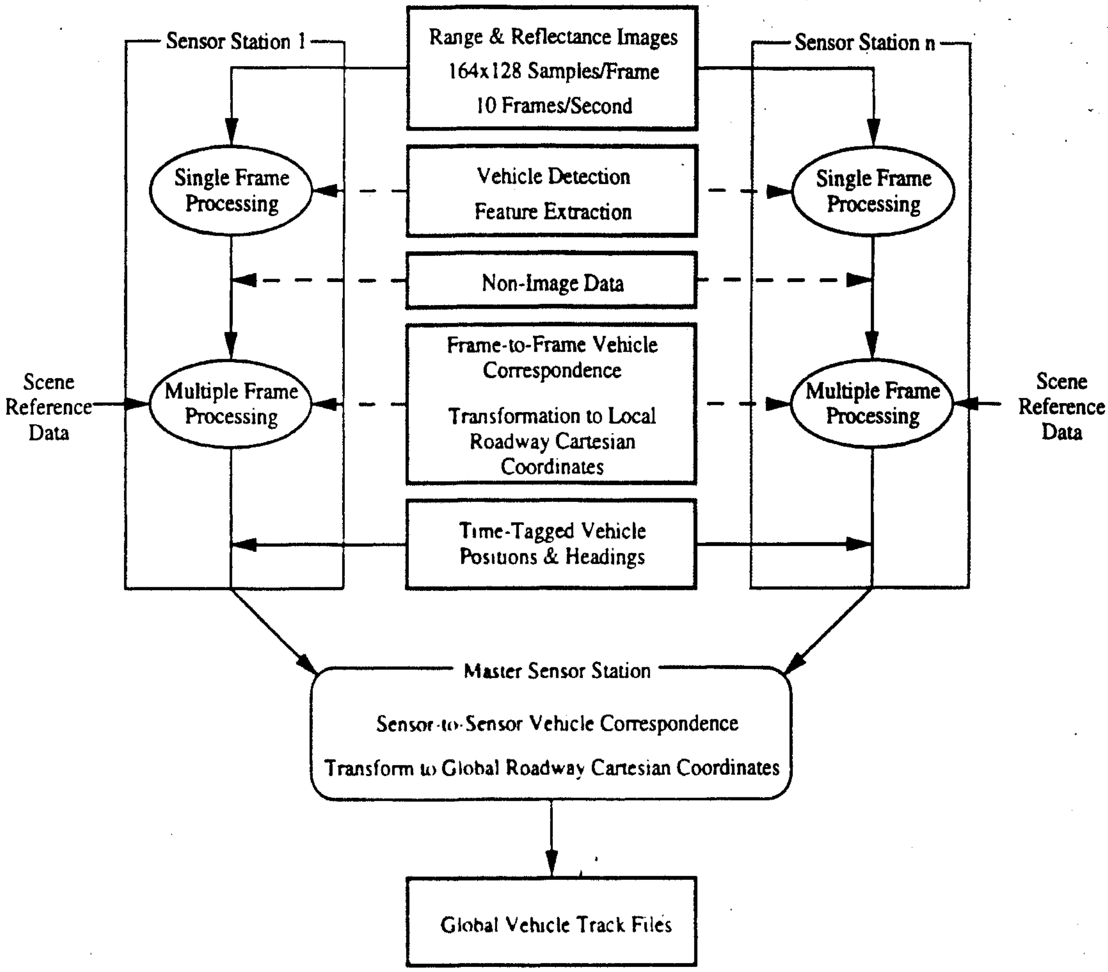

Figure 4-6. Real-Time Image/Data Processing

4.1.2 Sensor System Design and Implementation

The sensor system selected for the VME-MS is a modified, commercial, off-the-shelf laser radar manufactured by Perceptron of Farmington Hills, MI. It is comprised of a tower-mounted laser sensor head (LSH), a trailer-mounted laser sensor controller (LSC) and electrical interconnecting cables (Figure 4-7 a&b).

The 110-foot long electrical cables enable the tower-mounted LSH to communicate with the trailer-mounted LSC when the tower is extended to its nominal 100-foot elevation. The LSH scans a 35 degree vertical (only 32.2 degrees are used) by 50 degree horizontal field of view and provides both intensity and range images of all pixels (picture elements) within the subtended scene. When deployed under typical set-back conditions on its 100-foot high observation tower, a single Sensor Station can observe an approximately 200-foot long section of a typical five-lane roadway.

Laser Sensor Head

The LSH utilizes an amplitude modulated continuous wave (AMCW) direct detection scheme. In this approach, the voltage applied to a GaAs laser diode is modulated sinusoidally, in turn modulating the amplitude of the transmitted signal at the same frequency. The laser beam is scanned throughout the scene utilizing a polygon and nodding mirror arranged in tandem (the scanning pattern projected onto the roadway is illustrated in Figure 4-2). The reflected signal traverses the return optical path and is detected by a silicon avalanche photo diode and converted to an electrical signal. The phase shift of the return signal is measured relative to the phase of the transmitted signal, and therefore range is only measured unambiguously within a 15-meter interval defined by the product of the speed of light and the modulation period. Absolute range can easily be determined, however, because the nominal range to the road is approximately constant and known.

Sensor Mechanical Design

The LSH is an environmentally sealed unit that contains the scanning mechanism, laser transmitter, detector, and the associated optics and electronics. The transmitted and reflected laser energy passes through a flat glass window sealed to the enclosure. All electrical connections at the LSH are environmentally sealed. The scanning mechanism consists of a reflective polygon rotating in precision ball bearings and a planar nodding mirror. The polygon motion creates the 179 columns (only 164 are used) and the nodding mirror motion produces the 128 rows that together comprise the raster-scanned image. The LSH is mounted at three semi-kinematic bosses on its baseplate to a two-axis adjustable mount. The mount is adjustable over a 360 degree range of azimuth angles and over a +/- 10 degree range of elevation angles about its nominal 30 degree elevation position.

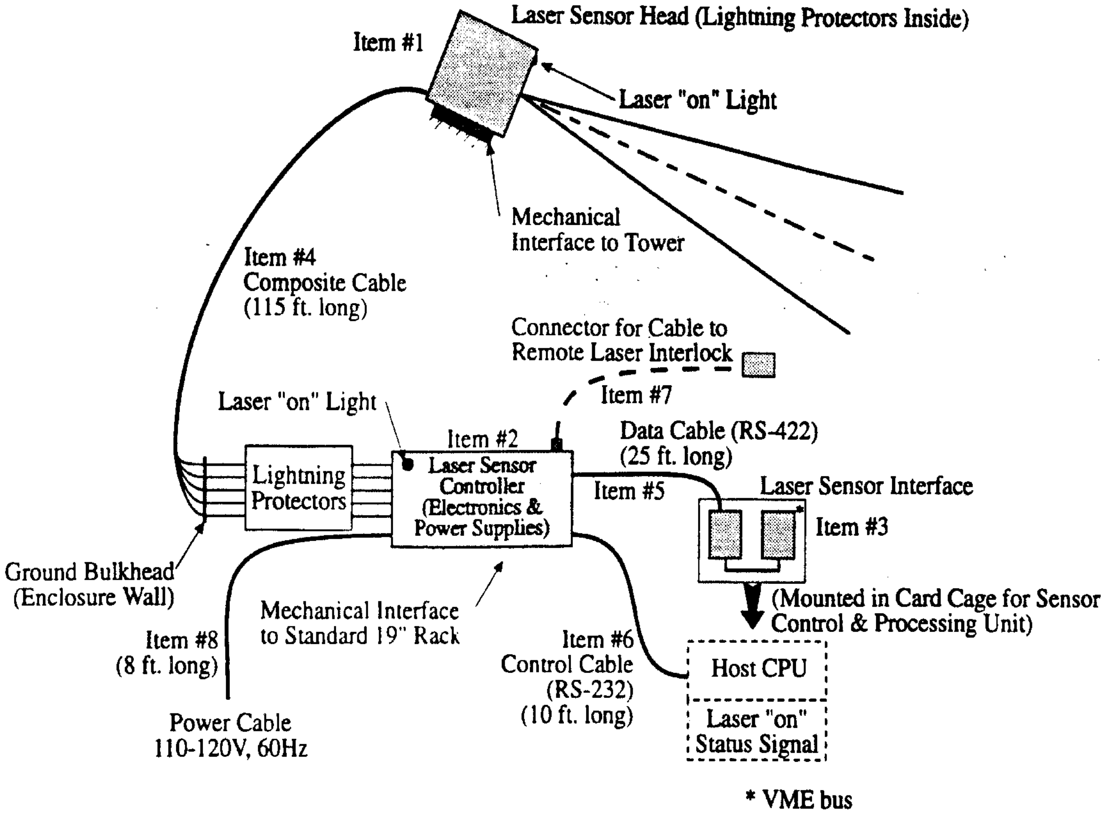

(a) Block Diagram

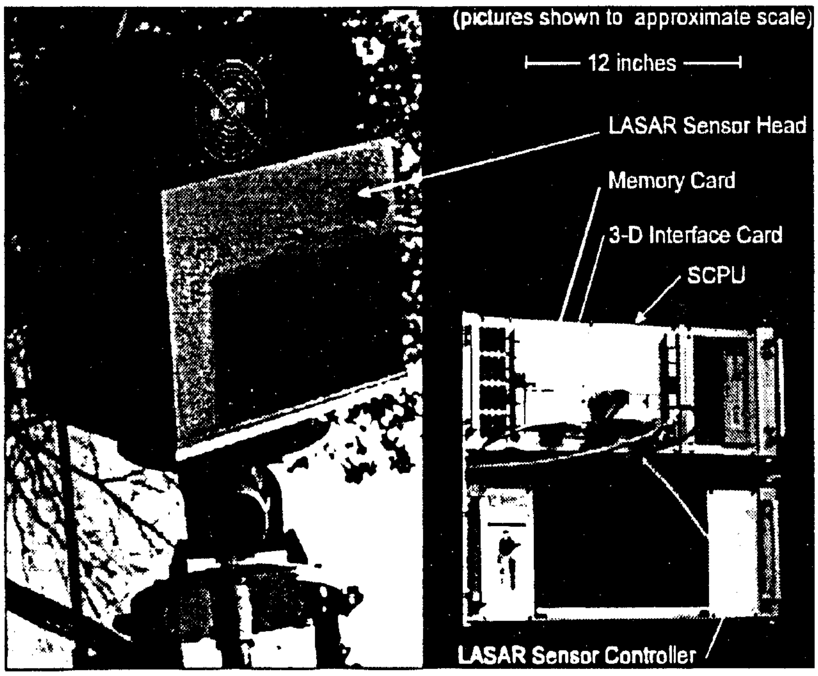

(b) Photos

Figure 4-7. Laser Sensor System

Sensor Optical Design

The LSH consists of transmitter and receiver optical channels. The transmitter channel consists of a GaAs diode laser and collimator assembly, a laser oscillator, and a laser bias amplifier. The laser diode operates in the near infrared spectral region (wavelength = 985 nanometers) and is amplitude modulated by the oscillator at 10 megahertz. The receiver channel consists of a narrow optical bandpass filter and collecting optics that focus the reflected laser energy onto an avalanche photo diode. The output of the detector is amplified and routed to the trailer-mounted electronics for digitization and processing.

Sensor Environmental Design

The LSH is provided with active cooling to dissipate the heat generated by the polygon scan motor and the laser diode assembly. An air-to-air heat exchanger is mounted to the top surface of the enclosure and is designed to remove heat from the air inside the enclosure. A thermoelectric cooler assembly is mounted to the rear surface of the enclosure and removes heat from the laser diode assembly via conduction through solid copper rods. A temperature sensor inside the enclosure reports the laser diode assembly temperature to a "watch-dog" control circuit located in the trailer-mounted electronics. If the upper set-point temperature of the LSH is reached, the polygon motor is shut down and the thermoelectric cooler is permitted to run until the low temperature set-point is achieved. Electrical surge protectors are provided in the LSH to shunt any lightning-induced current surges to ground.

Figure 4-8 shows a block diagram of a Sensor Station. This figure shows the Perceptron LASAR DatacameraTM Scanner (highlighted) and the Cohu video camera reside on top of the tower. A Perceptron-supplied cable connects the scanner to the Grounded Bulkhead Assembly (GBA), while a Cohu-supplied cable connects the video camera to the GBA. In the GBA, lightning protection equipment is utilized to arrest lightning current surges that may strike the tower.

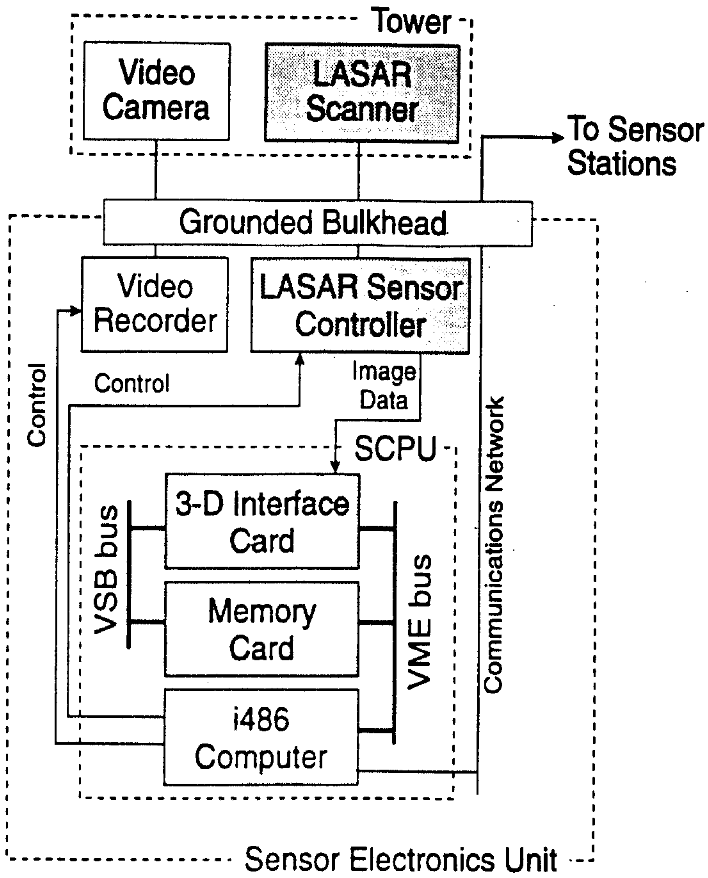

Figure 4-8. Block Diagram of a Sensor Station.

Sensor Electronics And Interfaces

The signals from the scanner are connected to Perceptron's LASAR Sensor Controller that controls the scanning mirrors and converts the range and intensity data of the scene into digital samples. The range and intensity digital data are delivered to the SCPU for processing via the Perceptron 3-D Interface and Memory Cards, which resides in the SCPU's chassis. System level control of the LASAR DatacameraTM is performed by the SCPU through an RS-232 interface.

The video signal from the video camera is connected to a commercial SVHS video recorder for archival recording. No interface to the SCPU is made other than an RS-232 command line to the video recorder to start and stop the recording of the video imagery. Power is supplied to the video camera through the GBA by a commercially available 12V DC power supply.

4.1.3 VME-MS Mechanical Design and Implementation

The VME-MS is an installation of typically three or more Sensor Stations Assemblies (SSA), or Sensor Stations, that collectively provide surveillance of an extended section of roadway or roadway intersection. System mobility requirements are satisfied by mounting collapsible towers on roadworthy dual-axis trailers that may be towed over moderately rugged off-road terrain for installation at remote sites. No heavy excavation equipment or poured concrete footings or anchors are required for installation of the Sensor Stations. Each SSA requires approximately 1kW of 230V, 60 Hz electrical power at its weatherproof marine-style connector located on the trailer-mounted electronics. Set-back distances from the street are dictated by property easements and the 65-foot minimum guying radius. Temporary fencing of the site is typically provided as a prudent safety measure. The VME-MS design approach specified the use of commercial off-the-shelf equipment whenever possible. Specifications for all VME-MS equipment can be found in Appendix B, VME-MS Design.

Sensor Station Assembly (SSA) Design

Each SSA is comprised of a trailer-tower assembly (TTA), a laser sensor head (LSH), a video camera assembly (VCA), an adjustable sensor mount, a sensor electronics unit (SEU), and lightning protection equipment (Figure 4-9). Collapsible towers, composed of 30-foot long telescoping sections, are mounted to the trailers at two pivot points. A manually operated winch is used to pivot the nested tower sections to their vertical orientation. A battery operated winch actuates an integrated pulley and cable assembly that lifts the tower sections to any desired height between 30-feet and 100-feet. The LSH and VCA are installed and properly oriented prior to tower tilt-up. All electrical connections to the tower-mounted equipment are implemented through a 110-foot long ruggedized cable assembly that is routed through the center of each tower section.

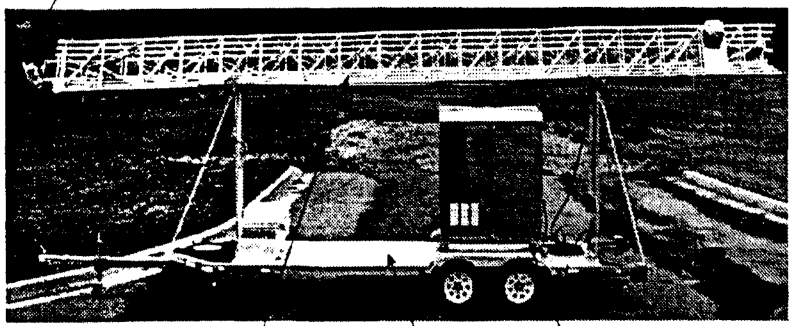

Figure 4-9. Sensor Station Assembly (Repeated)

All data processing, recording, housekeeping and communication equipment is housed in the weatherproof, double-wide SEU (Figure 4-10). Wire-rope vibration isolators, tuned to attenuate transportation-induced vibration, are installed between the SEU enclosure and the trailer deck All structural and fastening components are fabricated from corrosion resistant aluminum alloy, stainless steel or galvanized steel with the exception of the painted-steel tower outriggers. After arrival at the site, the 12 guy wires on each tower are anchored to manually-installed auger-style earth anchors. The earth anchors are proof-loaded to 2.5 times the maximum calculated guy loads based on a 75 mph wind condition.

Lightning Protection

Three levels of lightning protection are provided for each SSA. Level 1 is composed of an air terminal (lightning rod) mounted adjacent to the LSH, guy wire grounding kits and AC power line suppresser. Level 2 is comprised of grounding kits for the shields associated with all electrical cables. These shields are electrically bonded to the copper alloy grounding plate on the SEU. The grounding plate is connected to earth using wide conductive strips fabricated from copper sheet. Level 3 consists of a surge suppressor inside the MEU and LSH.

Electronics Packaging Design

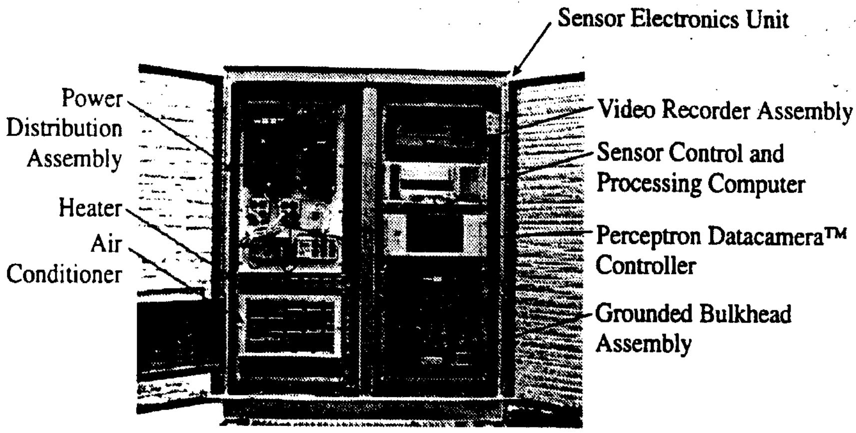

Figure 4-10. Sensor Electronics Unit (Repeated)

All sensitive electrical and electronic equipment is enclosed in the weatherproof, climate controlled SEU. The all-aluminum alloy SEU electronics enclosure, designed to be drip-proof when hosed down with moderately pressurized water, meets the National Electronics Manufacturers Association (NEMA) rating of 4X. All subsystems inside the SEU utilize standardized, modular, 19-inch rack-mount enclosures and mounting plates.

The SEU is composed of the electronics enclosure assembly (EEA), the laser sensor controller (LSC), the video recorder assembly (VRA), the sensor control and processing unit (SCPU), the grounded bulkhead assembly (GBA), the power distribution assembly (PDA), and the climate control system. The SEU is vibration isolated from its trailer and can be fully assembled, tested and then highway-transported without deleterious vibration effects. The temperature inside the SEU is thermostatically maintained between adjustable high and low temperature set points. Cooling is provided by a conventional 110V AC closed-cycle air conditioner. Heating is provided by an electric resistive-element heater.

The (PDA) is comprised of junction boxes with circuit breakers for separating the 230V AC into two separate 110V AC circuits that have balanced power requirements. The computers and other sensitive electronic components draw power from a "clean" 110V AC circuit equipped with additional electrical surge suppressors. The air conditioner and heater are connected to the second "dirty" 110V AC circuit.

4.1.4 Electrical Design and Implementation

The physical configuration of the major electrical elements is shown in Figure 4-11. The system consists of multiple Sensors Stations, each consisting of a standard video surveillance camera and 3-D laser sensor mounted on a tower, and a Sensor Electronics Unit at the base of the tower. Each station's SEU contains a Commercial off-the-shelf Video Recorder to archive the video data for later viewing and a Sensor Control and Processing Unit (SCPU). The SCPU controls the 3-D laser sensor's data acquisition and processes the 3-D image data.

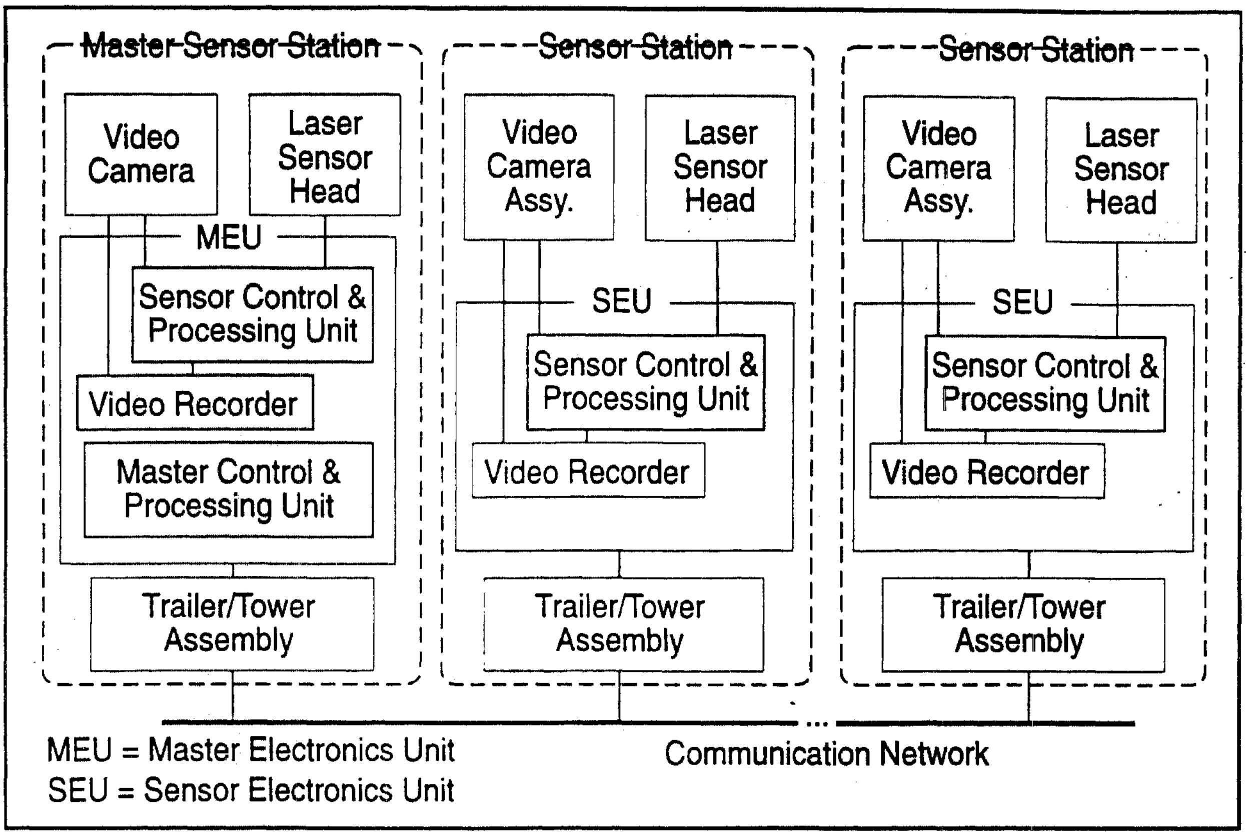

Figure 4-11. VME-MS Block Diagram (Repeated for Convenience)

The design details of the principal electrical elements of the VME-MS, and the Master and Sensor Electronics Units are now discussed.

Sensor Electronics Unit

Figure 4-12 shows the detailed layout of a Sensor Station with the SEU highlighted. This unit, which is housed in the enclosure on the Sensor Station's trailer, contains: a LASAR Sensor Controller, a Video Recorder Assembly, and a Sensor Control and Processing Unit (SCPU). Section 4.1.2 gives more information on the Sensor Controller.

Video Recorder Assembly (VRA)

The VRA is an commercial off-the-shelf time-lapse SVHS Video Recorder which archives the video data from the camera at 10 frames/sec, providing sufficient storage for recording 24 hours of continuous operation on one standard video cassette. The operation of this unit is controlled by the SCPU so that only video data corresponding to processed data is recorded. Specification 2498012 in Appendix B for more information on the VRA hardware.

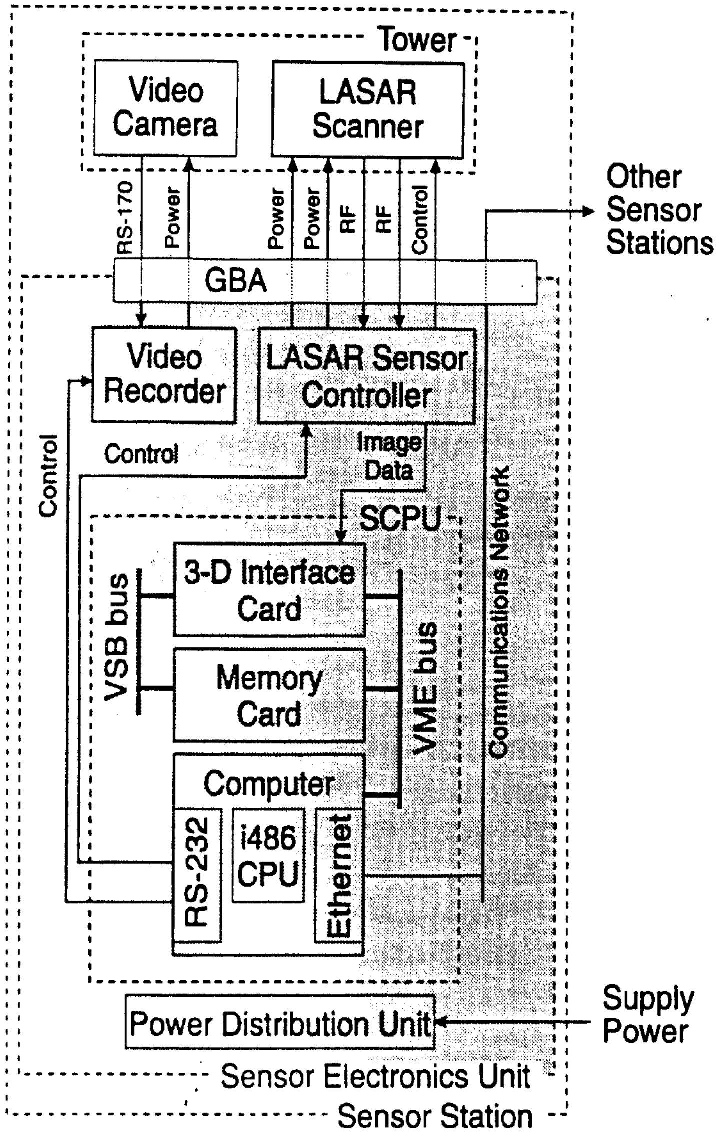

Figure 4-12. Sensor Station

Sensor Control and Processing Unit (SCPU)

The SCPU is a VME-bus backplane chassis which contains a i486DX2/66 single board computer (SBC), the Perceptron 3-D Interface Card, and a Perceptron-supplied dual-ported Memory Card which buffers the range and intensity data from Perceptron's interface card. The SBC accesses the range and intensity data from the Memory Card and processes the data (see discussion in next section under Image and Data Processing Software). Local vehicle track files are produced from the sensor data, and sent to the MCPU, which is housed in the Master Sensor Station. (See Specification 2498010 in Appendix B for more information on the SCPU hardware.) Figure 4-13 schematically illustrates the flow of data from the LASAR Scanner to data packets transmitted to the Master Control and Processing Unit over the Ethernet communications network.

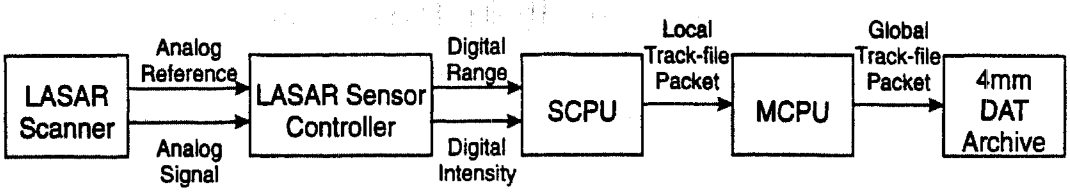

Figure 4-13. Data Flow Schematic

The SCPU controls the data collection and processes the LASAR Sensor data to generate a local track-file of vehicles in the sensor's field-of-regard. The SCPU also controls the VRA to record the corresponding video data during processing.

The Master Electronics Unit

The Master Sensor Station configuration, with the Master Control and Processing Unit (MCPU) highlighted, is shown in Figure 4-14. The Master Sensor Station is a standard Sensor Station with the addition of a Master Control and Processing Unit (MCPU) to a Sensor Electronics Unit converting it to an MEU. The MEU performs all the functions of an SEU plus the global functions of the VME-MS. (e.g.., high level control, global track file generation).

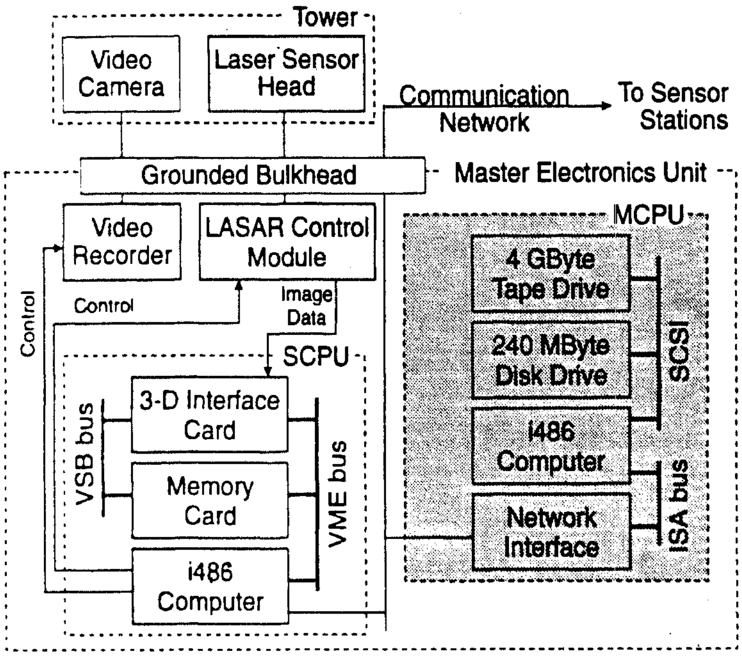

Figure 4-14. Master Control and Processing Unit

The MCPU is a commercial off-the-shelf i486DX2/66 industrial computer with a 4 GByte SCSI DAT Drive, a 240 MByte SCSI Hard Drive, and a 10Base-T Ethernet Network Interface Adapter.

The MCPU connects to all the Sensor Stations' SCPUs, including the one in the Master Sensor Station, via a 10Base-T Ethernet Local Area Network. The MCPU performs the synchronization of the sensor stations and merges the local track files of all the stations into one global track file. This track file is in turn archived on a 4 gigabyte DAT drive using the ANSI standard digital format for 4mm DAT media, which supplies a sufficient capacity for at least 24 hours of operation.

As well as performing those functions, the MCPU is a file server. Each SCPU SBC connects to the MCPU for its software, which is stored on the MCPU's hard disk drive, using Novell's Personal Netware network software. This creates a central software control model which facilitates SCPU software updates.

4.1.5 Software Design and Implementation

The Vehicle Motion Environment (VME) is a demanding real-time application which involves controlling several imaging sensors and processing their data to produce Global Track files. Figure 4-15 illustrates the relationship between the Master Control & Processing Unit (MCPU) and each Sensor Control & Processing Unit (SCPU). Each Sensor Station observes a roadway segment which is approximately 200 feet in length and 60 feet in width. As vehicles travel through that segment of the roadway, the task of the SCPU is to determine each vehicle's location, i.e., (x,y), and heading angle ten (10) times per second. Once a vehicle has left a Sensor Station's field-of-view, the local track file for that vehicle is sent to the MCPU over the local area network. The MCPU is responsible for tracking each vehicle on a global basis and merging the individual track files into a seamless global track file for each vehicle.

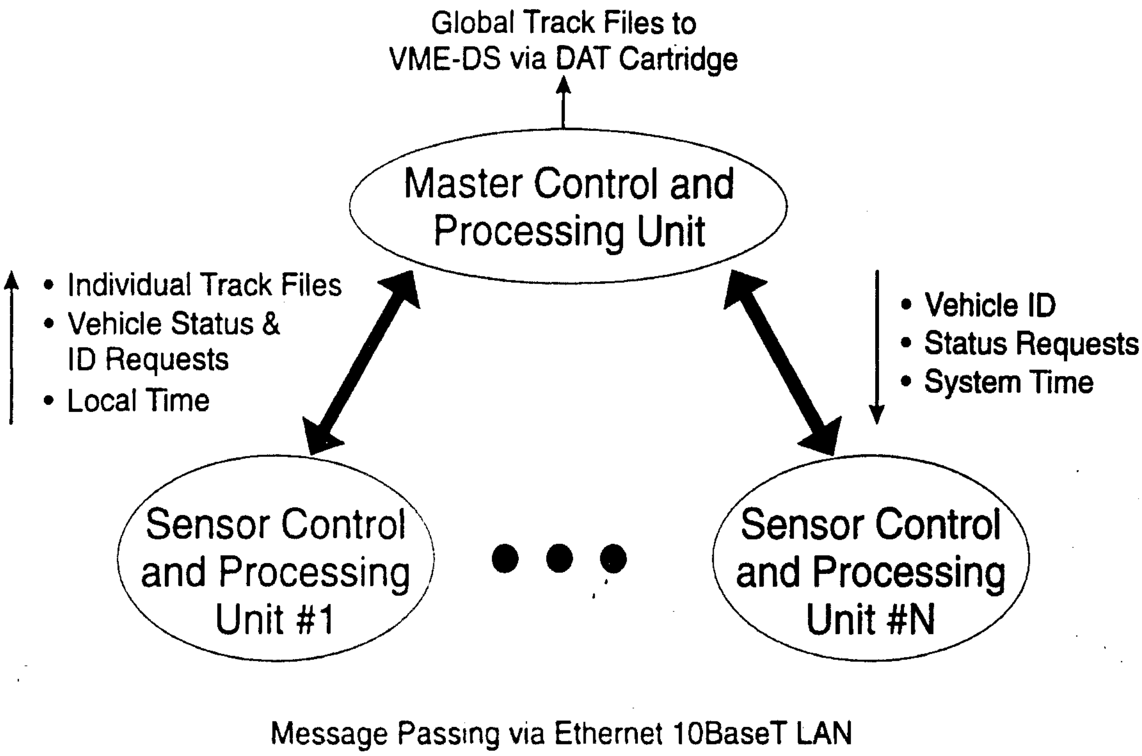

Figure 4-15. Relationship Between MCPU and Each SCPU (Figure 4-5 repeated for convenience)

Real-time applications are event-driven. Several processes on each computer must have sufficient and timely access to computing resources to respond to a stream of external events. For the VME project, the relevant time is determined by the fact that every 1/10th of a second each sensor begins collecting a new image of its section of roadway. Within these time intervals, the "events" are the number of vehicles that happen to be within a sensor's field-of-view. There are two major elements to the software design and implementation for the VME-MS: the operating system and the image processing.

Operating System Software

The issues that must be addressed in the selection of an operating system are: the image processing throughput; multiple sensor station coordination; and resource management and process management. Analysis of the image processing showed that a 66 megahertz 80486DX2 could handle the real-time data load. The computational burden of the MCPU is much less than the computational burden of a SCPU. This analysis assumed that the SCPU's would be operating in a "protected" 32-bit addressing mode. This is not the same environment in which most Intel computers reside. This fact eliminated most of the computer environments from consideration. We choose Poly FORTH from Forth, Inc. as our software development environment. This environment provides a 32-bit operating system. The Microsoft DOS is captured and "locked" into a "virtual" mode allowing the VME procedures to have access to the 16 bit DOS services to the file and network system. The assembler included with FORTH provides good access to the hardware, such as the sensor units, and provides the efficiencies required for the real-time image processing components. As with most FORTH systems, the extendible and flexible nature of the operating system can be made transparent to remainder of the software. The compact nature of the resulting software environment would permit an isolated SCPU to be tested in a floppy disk-only mode if the network failed. A standard ethernet-based local area network gives each of the SCPUs access to the hard disk drive on the MCPU.

FORTH is able to exploit the packet-based message passing protocol so that system components acquire a strong object-oriented nature. Software components are activated by the arrival of a message from another component. The message-based communication model manifest strongest on the MCPU. The MCPU is allowed to view the array of SCPUs as simply objects to which it can send and receive messages.

Table 4-1 lists the primary requirements that the VME Operating System had to meet. A 32-bit version of the FORTH operating system was selected because it was the only commercially available product which met all of these primary requirements. The need for the operating system to be extendible was driven by the fact that some operations had to run in 32-bit protected mode, while other operations had to run under the real mode operating system.

A high-end PC is used as the computing platform for a) the Master Control & Processing Unit (MCPU), b) each of the Sensor Control & Processing Units (SCPU), and c) the Diagnostic Computer. Figure 4-16 shows the internal computer architecture which serves all three computing environments.

Table 4-1. Requirements for the VME Operating System

(	Able to support Real - Time Operation
(	Less than 256 Kbytes in size
(	Extensible
(	Flexible
(	Able to be targeted onto a High-End PC
(	Transparent to the end users

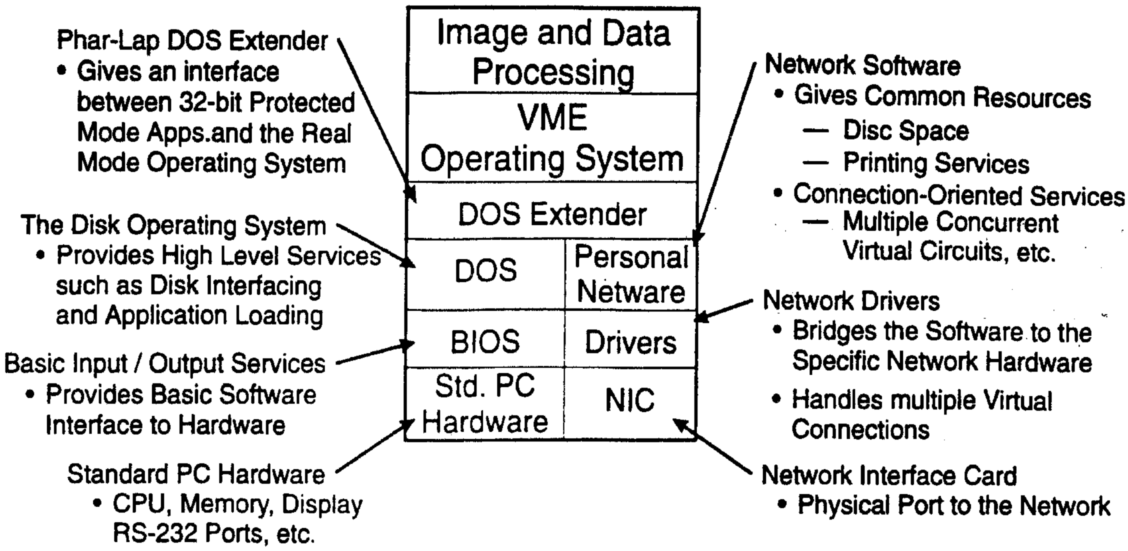

Figure 4-16. The Internal Computer Architecture

"FORTH" is a trademark of FORTH, Inc.

The end result was that the VME Operating System software provides the best of two (2) different computing environments: A 32-bit protected mode for conventional application software, and a real mode for meeting the unique needs for our very demanding real time application. The logic that constitutes the VME Operating System is presented in Figure 4-17.

The system configuration for the VME is illustrated in Figure 4-18. An Ethernet Local area network (LAN) is used to tie the individual Sensor Control & Processing Units (SCPU) to the Master Control and Processing Unit (MCPU). All communications between processors are handled by message passing using a standard protocol. The diagnostic computer is also tied in via the same LAN. From the viewpoint of a user, both the inter-processor messages and the track files produced by the VME system are completely standard.

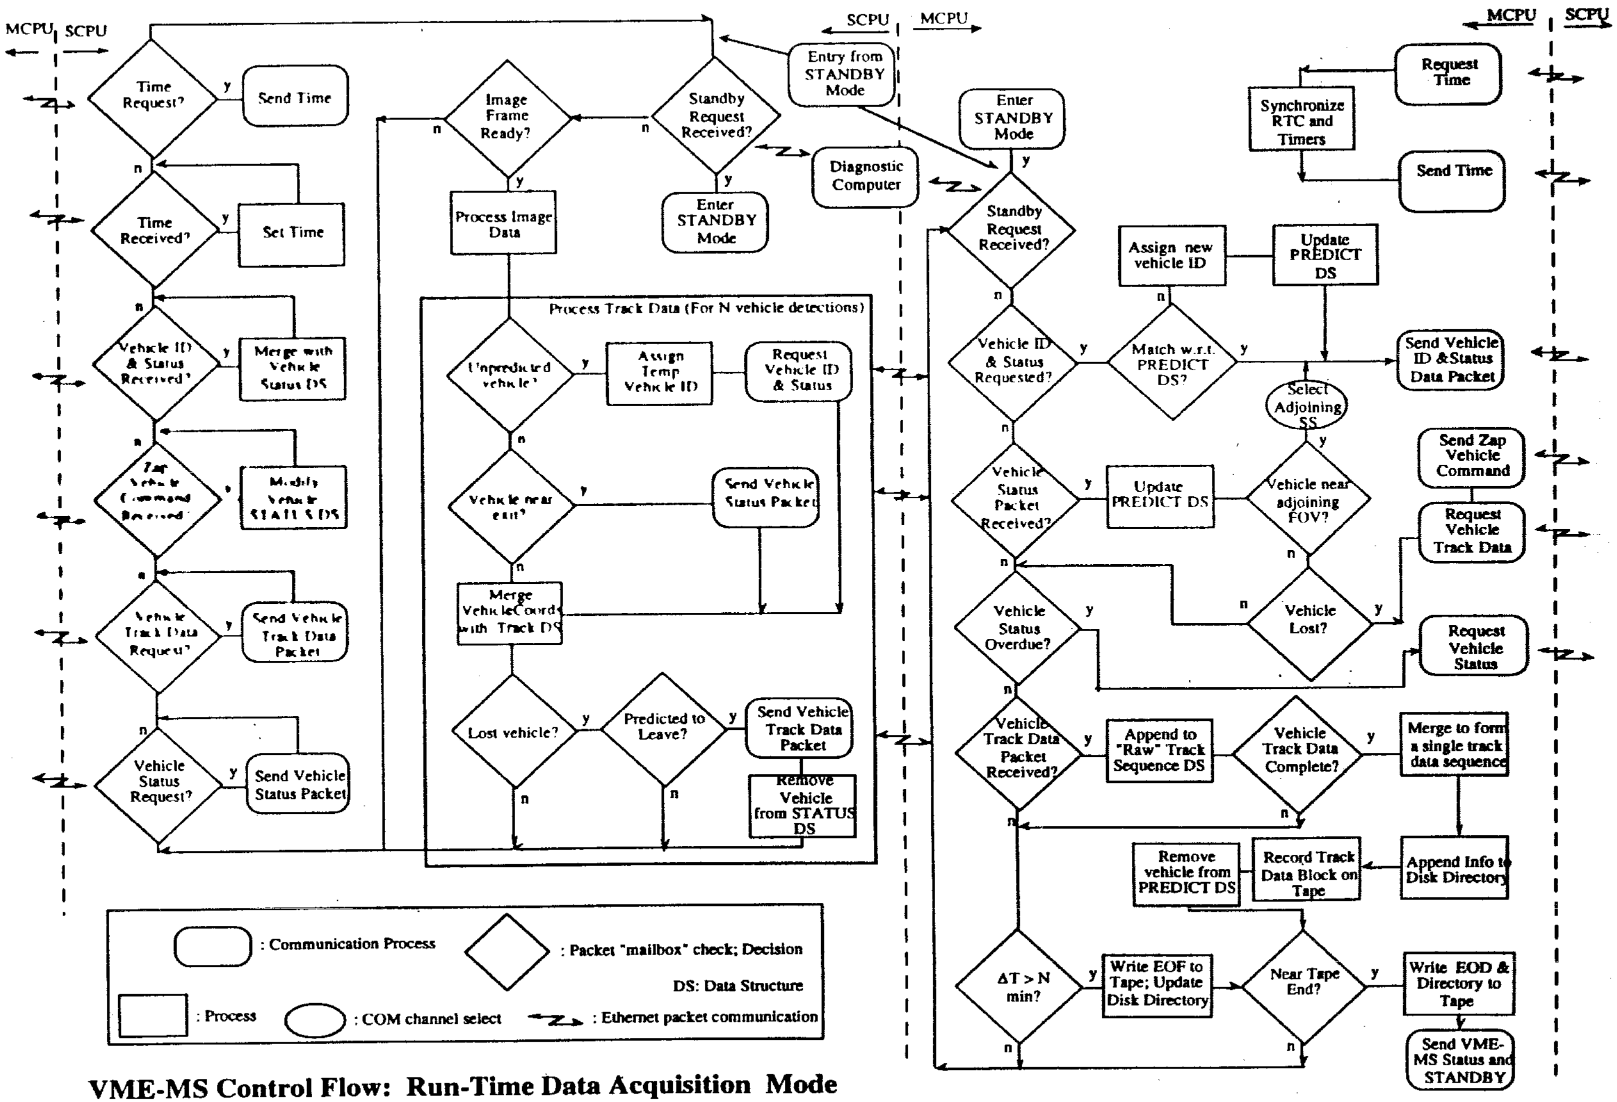

VME-MS Control Flow: Run-Time Data Acquisition Mode

Figure 4-17. Control Flow in Run-Time Data Acquisition Mode.

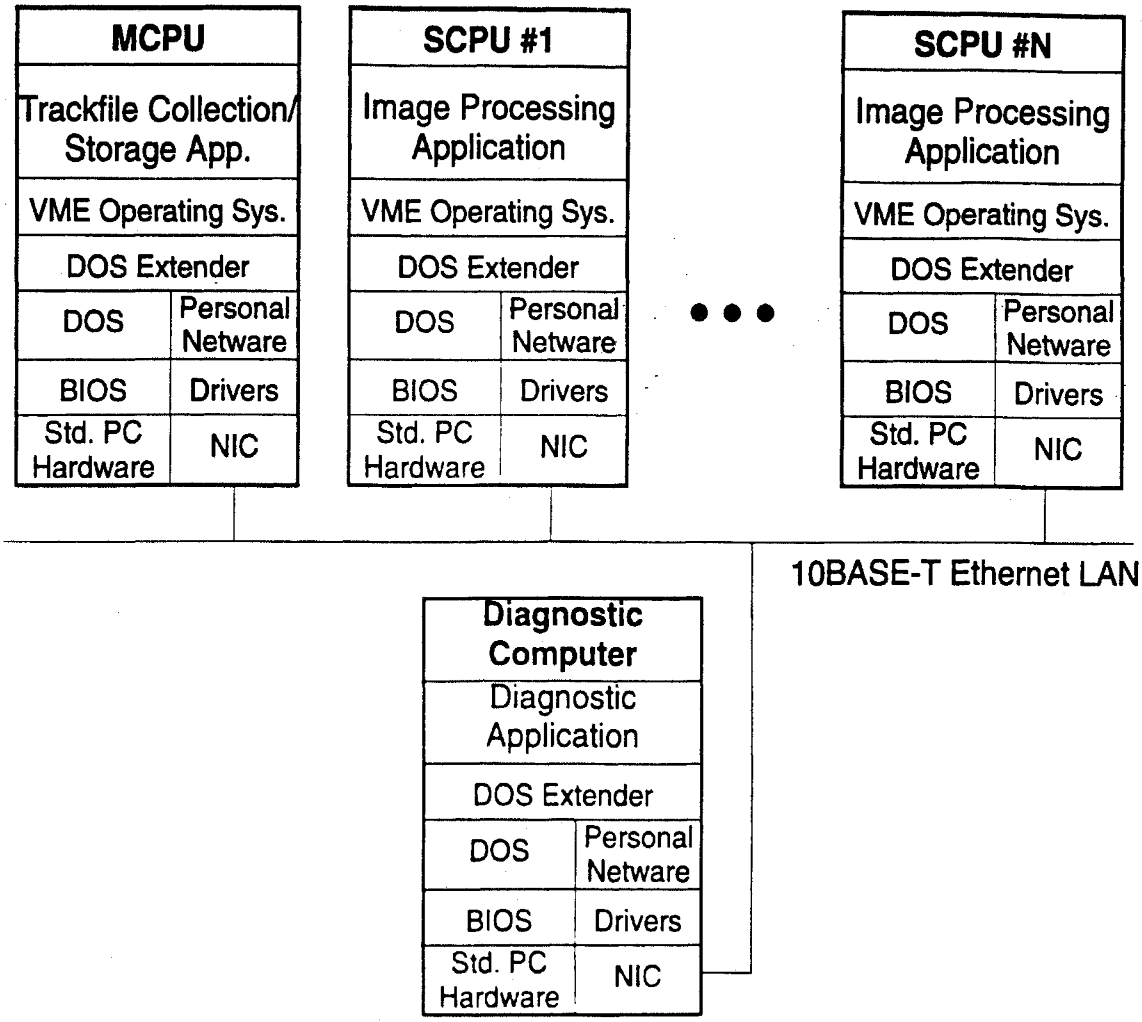

Figure 4-18. Configuration of the VME-MS System

Image and Data Processing Software

The VME measurement system is designed to observe a multiple, 200 foot section of a road system and produce "global vehicle track files." These track files contain observed vehicle statistics such as length, width, and position each time they are viewed by a sensor. From the user's point of view, the three individual sensors will appear as one sensor with an "exotic" field of view. The track files get written to magnetic media so they are available for later analysis by the user.

For this implementation there are four computers which act as a team to track vehicles across the combined fields of views of the three sensor array. The software design is based on using functional modules. Each module performs a clearly identified part of the job of producing track files. There are several guiding principles which help formulate the overall software system design. These include:

(1) The individual Sensor Station CPUs should execute identical programs. Any differences between their programs should be controlled by parameter files which contain the location of the sensor and other station specific parameters.

(2) There should be no direct communication between Sensor Station CPUs. This greatly simplifies the station specific information. If the Sensor Station CPUs were

allowed to communicate with one another, then each Sensor Station would need to know the field of view of the other stations.

(3) The choice of what task resides on which computer is based solely on the location of the required information and the available computer resources. For example, image processing tasks must be located on the Sensor Station that collects the data. It is possible that some task might be moved from the Sensor Stations to the Master Control Processing Unit in order to meet real-time constraints. Moving a task from a Sensor Station computer to the Master control computer requires analyzing the time it takes to transfer the required information from all of the stations, the time it takes for the master control computer to process this transferred information, and the impact of the additional data transfer on the local area network.

The software design for the information processing task is given below. The processing task is partitioned by which type of computer (master computer or sensor station) does the task.

Data Processing on the Master Control and Processing Unit (MCPU)

The MCPU has several tasks. Most of these tasks are associated with the production of track files, such as writing a completed track file to tape. Other tasks are management tasks which insure that the station computers get the information they require. The Sensor Station-to-Sensor Station vehicle correspondence is completely dependent on the MCPU collecting and forwarding vehicle tracking information to each Sensor Station. The best way to understand the role of the MCPU is to examine the information and hardware environment of this computer.

The hardware resources are simple. The MCPU is the only one with a real hard disk. The SCPUs must access this hard disk over the network via Netware. The most important device as far as processing track files is concerned is the magnetic tape drive. Information resources include:

(I1) The location of the field of view of each Sensor Station (see figure 4.5.2.1)

(I2) Road and lane locations

(I3) Local track files from each sensor station

(I4) Vehicle track parameters which contain position, time and velocity information. These track parameters are sent to the MCPU from a Sensor Station just before the vehicle enters another station's field of view (FOV). The MCPU must determine which (if any) FOV the vehicle will enter next. The track information must be sent to the appropriate Sensor Station CPU before the vehicle enters its FOV. This is the mechanism for solving the vehicle correspondence problem between Sensor Stations.

The two most important functional components of the MCPU are described here.

The **Global Track File Manager (GTFM)** is responsible for receiving the local track files from each Sensor Station and aggregating them into one seamless track file.

There are two types of information in the track files. The first type is the vehicle descriptors such as length, width, and master ID. The other type of information is time-varying data such as location, time of observation, and the observing station. On a vehicle-by-vehicle basis, the GTFM must combine corresponding vehicle descriptors from several stations into one. The vehicle position entries in the track file must be sorted by time. After both of these things have been done, the global track file is written to tape and the global track file for that vehicle is closed.

The **Prediction** component on the MCPU has several functions. These include:

P1: Assigns a global ID to a vehicle when a station requests it. This ID stays associated with a particular vehicle while the vehicle is in the combined field of view.

P2: Maintains velocity and position estimates for the vehicle. When the ID is assigned the velocity estimate is set to the average velocity for recent vehicles in this same collection of traffic. This information gets updated when **prediction** receives the vehicle track parameters.

P3: Sends a request for the vehicle track parameters if more than a user-defined amount of time goes by without receiving vehicle information. This task is checking for lost vehicles and other errors.

P4: Sends a command to a sensor station to send a vehicle track file and close out the track file. The track file will have an error flag with it. This command is sent only when too much time has elapsed without a vehicle report and the Master CPU have given up on the vehicle. The track file will be processed and saved with a "Lost the vehicle flag."

P5: Relay track parameters. The component is named for the task it performs. This component receives a track parameter file for a vehicle. It uses the position and velocity estimates of the vehicle to predict which sensors' fields-of-view the vehicle will enter next. The vehicle track file is sent to all possible sensors stations that might see it. This allows the next Sensor Station to open up a track file several tenths of a second before the vehicle is detected. This process permits a vehicle to be reliably tracked across several field of view of several Sensor Stations.

Data and Image Processing on the Sensor Control and Processing Units (SCPUs)

Each Sensor Station has access to a rich information environment. A new 164 row by 128 column image arrives 10 times a second. This environment is the 209,920 pixels (i.e., samples) per second of image data. Each pixel has two channels: range and intensity. Vehicles moving along the roadway must be detected using this image data. A detected vehicle must either be associated with an existing vehicle so that the appropriate track file can get updated or be declared a "new entry" for initializing the formation of a new track file. When the vehicle crosses one of several lines the vehicle's track parameters are sent to the MCPU. Finally, after the vehicle leaves this Sensor Station's field of view the track file is sent to the MCPU. The most important functional components are described below.

Range Reference Image Tracker (RRIT) maintains the range to the road when no vehicles are present and provides the basis for detecting the change produced by vehicles . Two methods are used to obtain the range to the road when cars are not present or rarely present. The primary method is to use the vehicle detection image to choose which pixels to update. This allows only road pixels to have their range updated. There is a small problem with this approach. The vehicle detection image is computed from the range reference image. But how does the range reference image get initialized? We elected to also use a second technique wherein the range to the road is estimated by a 75th percentile tracker. This procedure will correctly estimate the range to the road for a pixel when that pixel is not blocked by vehicles more than 75 percent of the time. As a extra measure of insurance, this initialization should only take place when the traffic is light.

There are two reference images. The first is a relative range reference and the second is an absolute range image to the road. The relative range reference image is used in vehicle detection. The absolute range reference image is used in computing the (x,y,z) coordinates of vehicles. The first image contains the modular range to the road. It has an error of 50, 100, 150, 200, 250, or 300 feet plus an inch or so. The second reference image is obtained from the first. It uses the assumption that there are no range discontinuities close to the above values. When a range discontinuity is detected, the system tries to add a multiple of 50 feet (really 49 feet 2.55 inches) to the modular range value. The procedure starts with the road closest to the tower and works its way away from the tower. The approximate range to the closest pixel is known prior to within a few feet from the knowledge of the tower and the road geometry. This allows the software to compute the number of multiples of 50 feet to the relative range estimate to obtain the absolute range. This gives the real range to a road pixel. The remaining steps assume that the range between adjacent pixels is less than 1/2 of 50 feet. The multiples of 50 feet to add to the relative range is that integer which minimizes the range difference to the adjacent pixel with a known range.

The **Vehicle Detection** component applies a simple threshold on the difference between the relative range reference image and the current range channel image. If there is a vehicle blocking sight of the road, then the range of the current image will be smaller than the corresponding range in the reference image. Pixels that are one foot, or more, closer to the tower than the road are flagged for further processing. This processing step is the last one performed on the entire 164 row by 128 column image. Detected pixels which form a connected region, called blobs, are analyzed by the next processing step. This step begins by rejecting all blobs that are too small to be a vehicle. The closest edges of the blobs form the initial feature set for the vehicle. The (x,y,z) coordinates for the blob are extracted, starting with the two closest edges and working into the blob until the (x,y) position of the new pixels get more than a foot away from the first edge estimate. The idea is to capture all of the pixels from vertical vehicle edges found on trucks. Pixels from the top of vehicles are not converted to (x,y,z) coordinates. The result of this process is a list of (x,y,z) points which are associated with the two closest edges of a vehicle. These points are partitioned into linear segments which are used to estimate the length and width of the vehicle. The corner of the vehicle is used as a temporary position estimate. (Later, this will be converted to the vehicle centroid.)

A newly detected vehicle is a vehicle detection not associated with any existing track file. The prediction component estimates where vehicles should be in this image. If a detected vehicle is where it is predicted to be, then the detected vehicle is used to update the appropriate track file. When a vehicle does not match anything in the prediction list then it is either a new vehicle or it is an error. In both cases, a new ID is requested. If the vehicle is detected in an area where new vehicles should be detected then no error flag is set. Otherwise a "Vehicle appeared in unexpected location" flag is set in the track file. The end user will have to address with this problem during the initial steps of data reduction. A different type of error occurs when a vehicle that is in the prediction list can not be associated with a detected vehicle. The position of such a vehicle will be updated for several frames until it is matched, or until a given amount of time passes or the master computer gives instructions to close this track file. In the latter case, this track file will have its "Vehicle lost" error flag set.

The **Prediction** component is the heart of the vehicle correspondence software, both within and between fields of view. This routine exploits the fact that during a tenth of a second vehicles cannot do much more than move linearly. The prediction component maintains a velocity estimate for each vehicle. When the vehicle is first detected, the prediction component sets the vehicle's velocity to the mean for recent vehicles in the lane in which the vehicle was detected. After the vehicle is detected in a second frame, the change in position gives a velocity estimate. When a vehicle is arriving from another field of view, prediction is given a position and velocity estimates for the vehicle just before vehicle should arrive. When the position of a vehicle crosses one of several lines, the vehicle's ID, position, and velocity estimates are sent to the MCPU.

4.2 VME - DS Software

This section of the report describes the VME data processing system (VME-DS). The VME-DS processes libraries of VME track files and allows users to view data in a variety of formats (text files, graphs, animations). The VME-DS also supports certain types of processing calculations on the raw track file data (smoothing, Kalman filtering, etc.). Special calculations are also included to allow users to conduct crash detection / near-miss calculations, or to export files of more detailed information on each vehicle's motion experience (e.g., range and angle-of-attack data for each vehicle pair in the field of regard). These latter data can then be further processed and analyzed with commercially available statistical analysis programs or other analysis tools to obtain histograms and other specialized plots describing the motion experiences of the selected files. Additional specialized calculation modules may be added subsequently depending upon anticipated needs. (For additional details regarding the material described in this section, the reader is also referred to Appendix C which contains additional documentation and user manuals for the current version of the VME-DS software.)

4.2.1 Basic Elements of the VME-DS Software

Figure 4-19 shows a diagram containing the principal pieces of the VME-DS program. At the top of the figure, raw track file data obtained from ERIM are initially pre-processed to regularize the data and to remove any obvious anomalies that can be readily detected at this stage. The data that emerge from this pre-processing stage are then assembled into a VME database. This database(s) may be quite large or very small depending upon the amount of data being pre-processed from ERIM (e.g., a few hours of track file data or a few days). The size of such database files can range from a few megabytes up to hundreds of megabytes. A number of such database files may then be assembled together as a total VME database corresponding to a particular road site and time period (e.g., seven data base files corresponding to each day of the week, together covering an entire week).

Small catalog, or index, files are created for each database file. A catalog file contains header-record information for all track files in its corresponding database. This permits the VME software to search for information more rapidly while also interacting with the user more efficiently.

The primary interface to the program user is the file manager and search module elements seen in the center of Figure 4-19. Having opened a particular database to process, the program user would normally select a time frame of available data to process. This may simply be a start time and an end time (e.g., all files between 7:00 AM Tuesday to 6:00 PM Wednesday). These data files could then be further sub-selected using the search module to obtain just truck-type vehicles (e.g., vehicles having lengths greater than 20 feet). The resulting subset or list of data may then be examined individually, or together, by performing various operations on these track files.

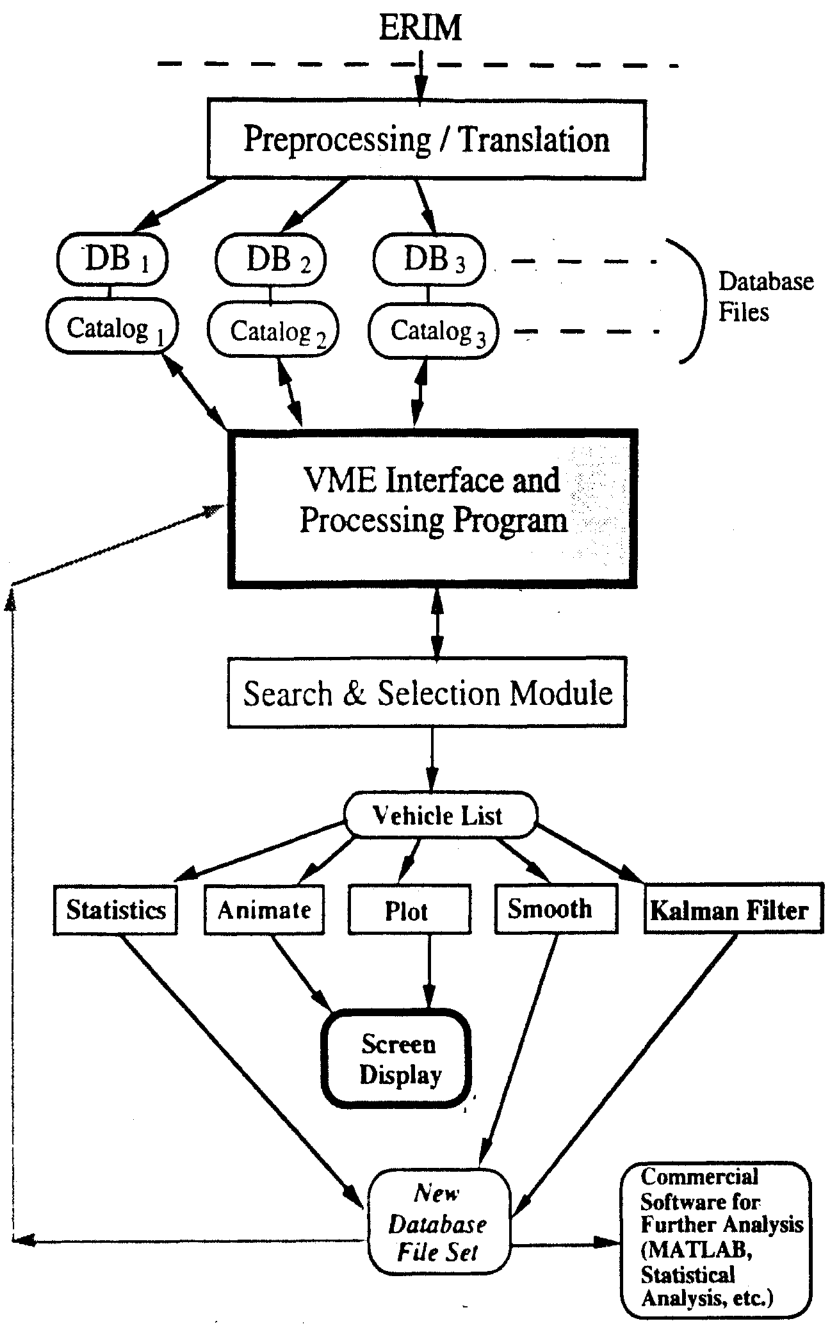

Figure 4-19. Basic Concept of VME Software & Processing Flow.

In the lower portion of Figure 4-19, several processing blocks are seen (animate, plot, etc.). Each of these blocks represents a possible action that the program user can select to process data from the vehicle list of available track files. A user might wish to first view a few individual files as text files, or to graph some of the track file data to verify its content. The user can also elect to animate all the available files as an on-screen playback of traffic flow. Alternately, the user may wish to conduct some preliminary smoothing of track file data and then an animation.

At a more advanced level, the user might elect to extract additional information from the raw track files by performing so-called Kalman filtering calculations. Kalman filtering is a technique that combines measured data with approximate models of the system being measured (cars, trucks, buses, etc.) in order to enhance the accuracy of the direct measurements. Importantly, this filtering process also permits estimates of additional system responses to be obtained. For example, the VME track files contain no direct measurement of vehicle longitudinal acceleration or driver steering activity — only vehicle position information (x, y, heading angle). By employing Kalman filtering calculations, additional information (that are implicitly contained within the positional measurement data of the raw track files) can be estimated and extracted from the positional measurement data. Thus, items like longitudinal acceleration and driver steering activity can then also be estimated using the Kalman filter option. In addition, the accuracies of the direct measurement positional data (x, y, heading) are also further improved as part of this filtering process.

Since Kalman filtering computations require more intensive numerical processing than other more routine types of smoothing calculations, not all files would ordinarily be selected for Kalman filter processing as a routine matter. However, if more detailed track file analyses are required, the built-in Kalman filtering module present in the VME-DS software can be used in this capacity. Further example discussions of Kalman filtering applications appear subsequently in this section and in Appendix C.

4.2.2 Two Modes of Operation

The diagram in Figure 4-20 shows the two basic modes for operating the VME-DS software. The first mode is an *interactive* mode in which the user would normally be examining smaller groups of files. This would normally involve such activities as looking at the contents of specific files using the text editor or graphical plotter. It may also involve animating short sequences of traffic flow to view the dynamic interactions of certain groups of vehicles in the traffic stream. Selected files may also be further analyzed in this mode using the Kalman filtering module to extract additional response variables for further review.

The other primary method of operating the VME-DS is in a *batch* mode normally used for processing large numbers of track files. Any specialized calculations that a user may wish to execute in a one-shot fashion on large groups of track files would normally be included under this category. Two such calculation options are currently built into the VME-DS software. One is for detecting crash incidents. The other specialized calculation computes inter-vehicular ranges and angles of attack for all vehicles in the traffic stream at each sample time. These calculation results are normally exported to a disk file for further post-processing and analysis using other software packages. Other types of special calculation modules can be added under this batch mode category in the future.

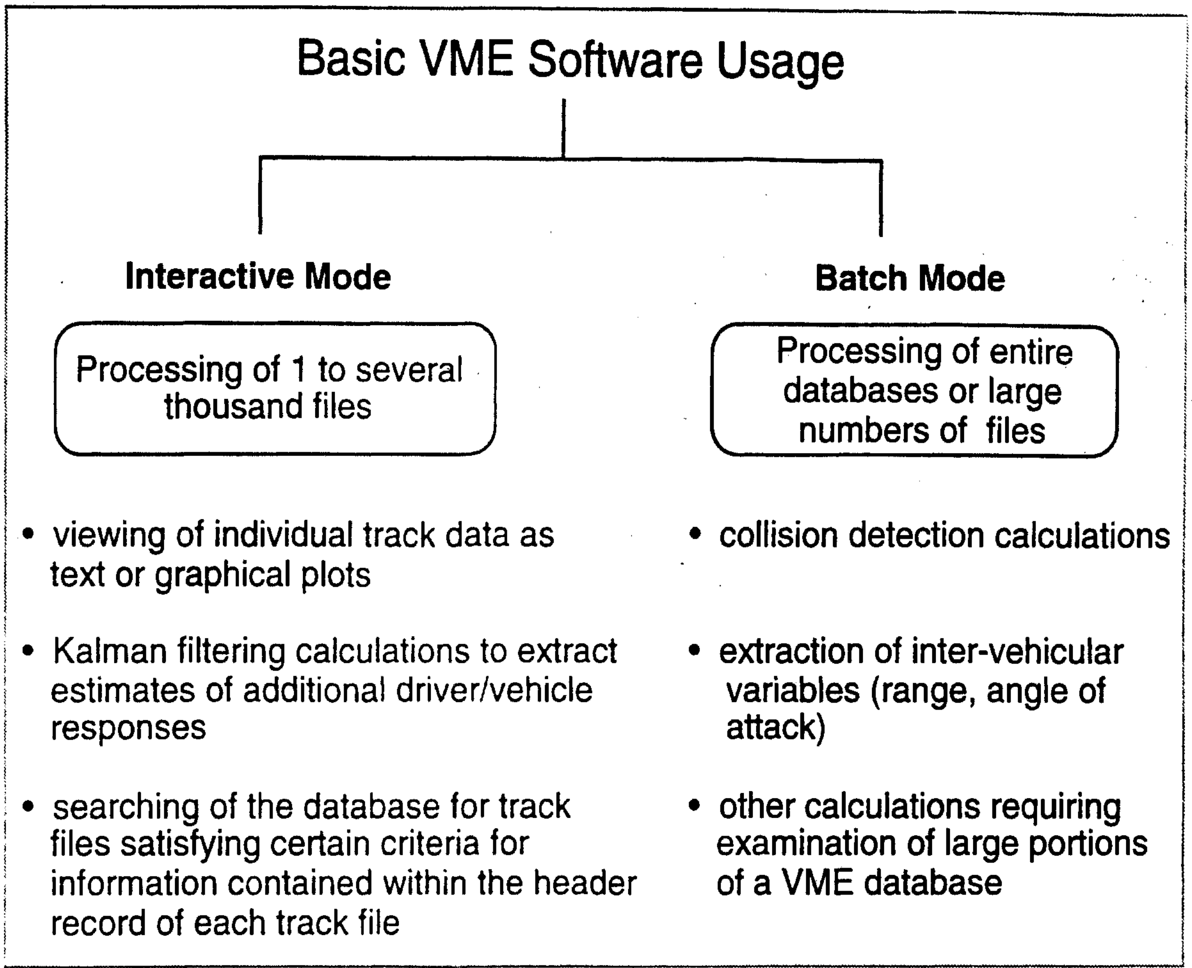

Figure 4-20. VME-DS Software - Modes of Operation.

4.2.3 General Architecture and Program Components

The VME-DS software is written in C++ and runs on Macintosh Quadra and PowerPC Macintoshes. Porting of the C++ code to other windows-based machines is not expected to be a major obstacle once the current code is finalized. Figures 4-21 and 4-22 show diagrams illustrating basic features of the preprocessor and VME-DS software in its present form. These diagrams describe the logical design of the software systems.

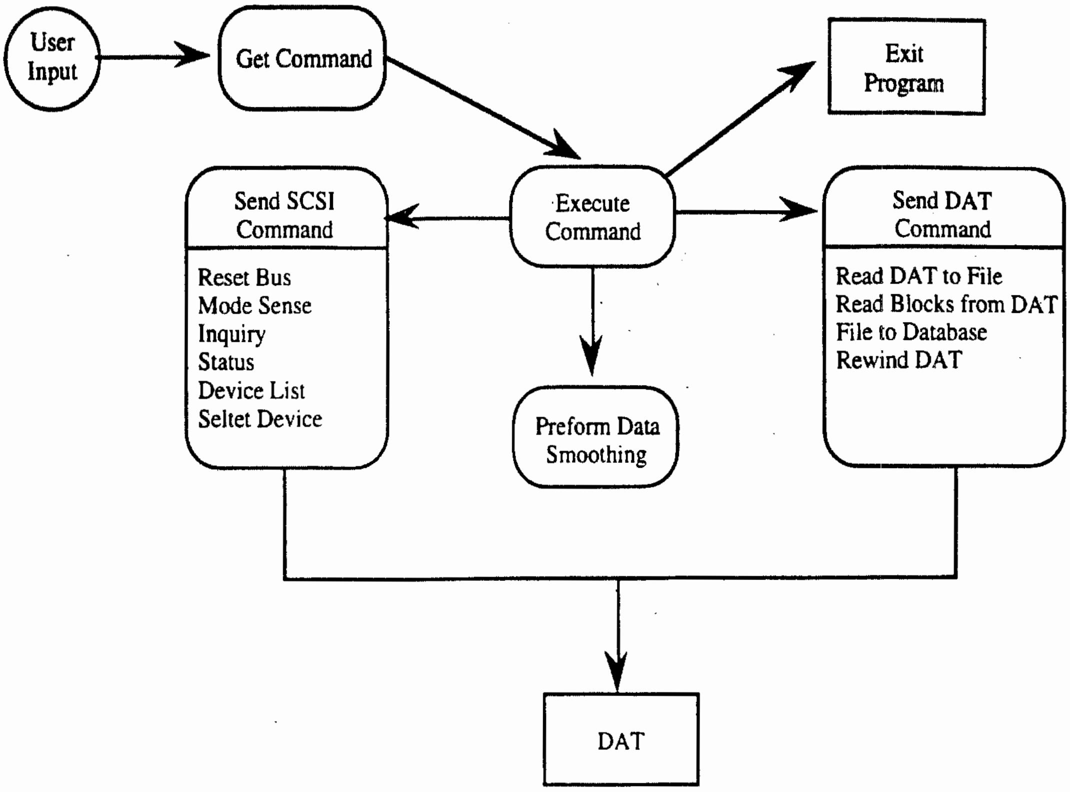

Figure 4-21. Pre-Processor Software.

Pre-Processor

The pre-processing element described in Figure 4-21 reads the DAT tapes supplied by ERIM and converts the ERIM file format to a Mac-specific file format. It also builds a VME-DS database, and creates a catalog/index file for each database. The preprocessor can also perform certain preliminary smoothing operations on the track file data and screen data for any apparent anomalies. For example, in cases where the ERIM sampling rate maybe be variable, the pre-processor would regularize the data through interpolation and re-sampling operations. The output product of the preprocessor is a VME database of track files for use by the VME-DS software.

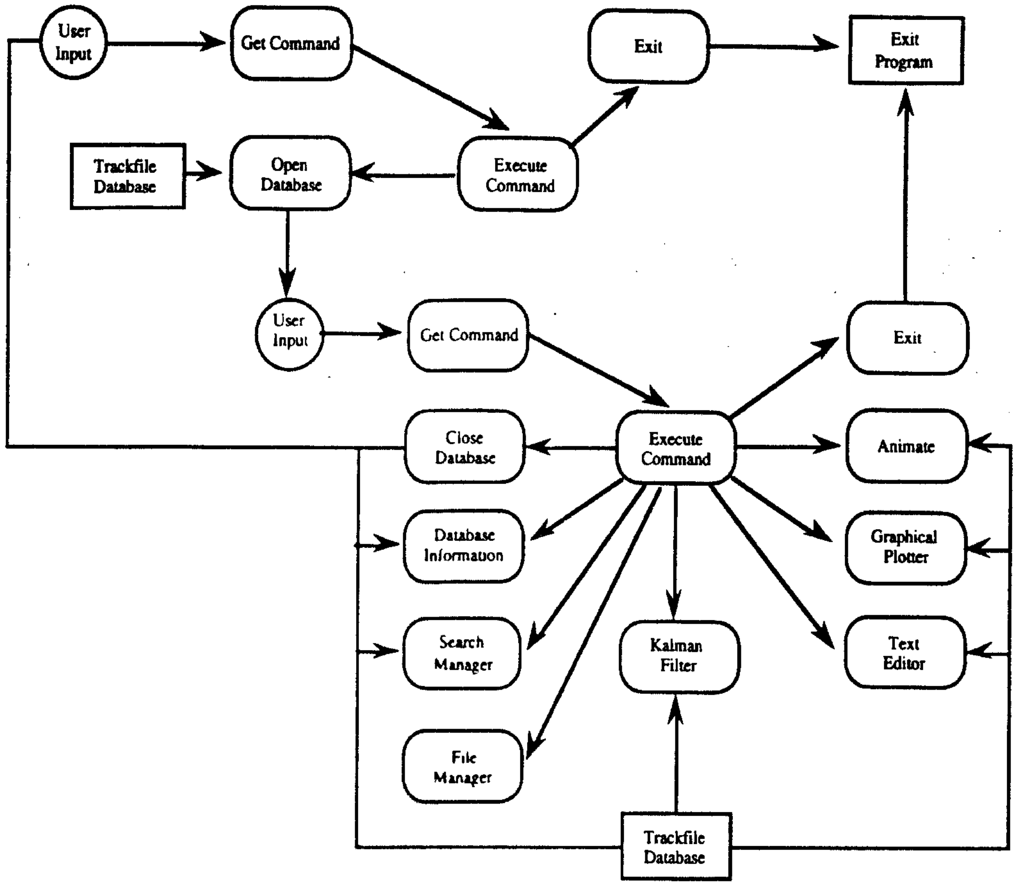

Figure 4-22. VME-DS Block Diagram.

Track File Database

A track file database is comprised of two parts - a site header record and M track files. The site header contains information specific to the site from which the data was gathered, such as the time of data collection, the total number of M track file records in the database, and the site geometry. This information is used by the VME-DS program in subsequent data processing operations.

A track file itself is comprised of two sections — a header record and N data records. The header record includes simple codes for identifying the type of vehicle, various status flags, for noting anomalies in the data, the time and speed the vehicle entered the field of regard, the length and width of the vehicle, and the number of data samples contained in the track file.

Data records describe the movement of the vehicle through the field of regard as time history samples. Each record includes the encounter time, the x and y locations of the vehicle centroid, and the yaw angle of the vehicle. An estimate of standard error for each measurement sample also appears in the data record.

Catalog/Index File

A catalog file is created for each database file and used by the database to facilitate searching and interacting with the program user. The file format is basically the same as the database format, absent the time history data records.

Search Manager

The search manager is used to query the database using a specific search criterion. The search manager supports two types of searches — primary searches and sub-searches. A primary search is defined as a search of the database that is not based on a previous search. A sub-search is defined as a search which is operates on the results of a prior search. Sub-searching lets a user begin with an initial search, such as time boundaries, and then further narrow the scope of the search using constraints on such items as vehicle length or other status flags contained in the header record.

The search manager contains a graphical browser window that displays the result of each search. Detailed search information is displayed for each search conducted by the browser. See Figure 4-23.

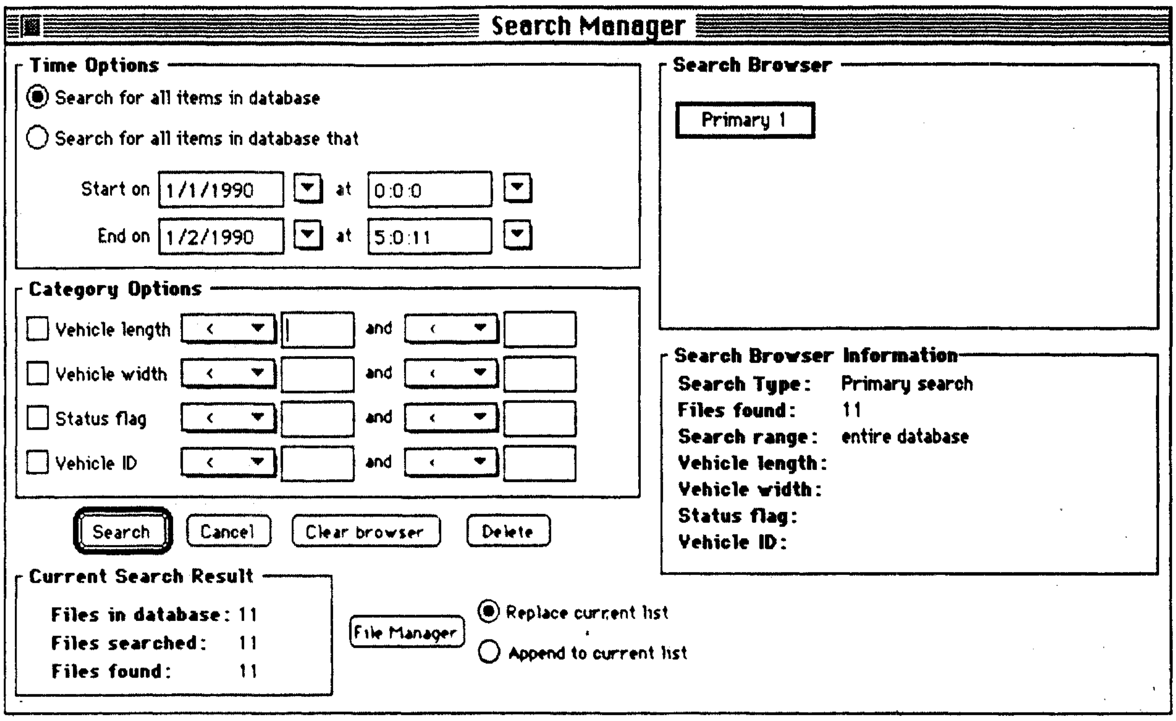

Figure 4-23. Search Manager Dialog.

File Manager

The file manager collects search results and allows the user to view and further analyze the results. Search results may be plotted, animated, or viewed in a text editor. In addition, the results may be Kalman filtered and viewed using the same analysis tools. An example is seen in Figure 4-24.

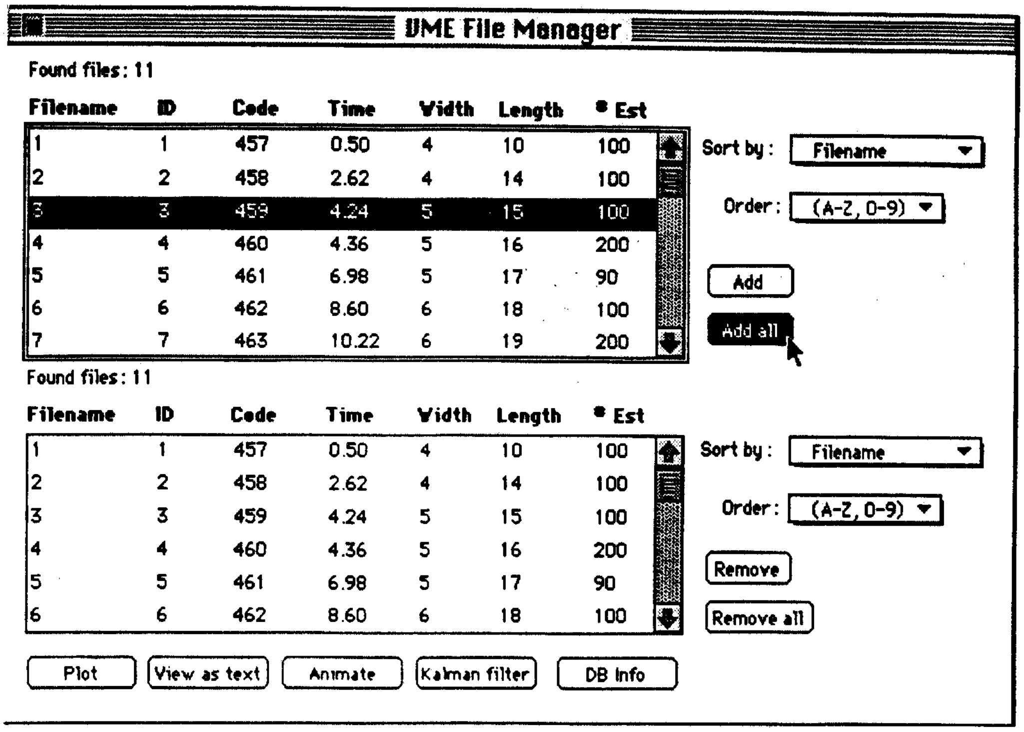

Figure 4-24. File Manager Dialog.

Text Editor

The text editor component permits the user to view one or more track files as an ordinary text file comprised of the header information and the track file measurements listed as time histories. Figure 4-25 shows such a listing with the header information listed at the top of the window. Samples of time history measurements then follow. The first column is Time and represents the time of passage of the vehicle from its entry point into the field of regard. The X, Y, and Yaw columns represent the sensor measurements of vehicle (centroid) longitudinal position, lateral position , and heading angle. The Error column will contain an estimate of positional error provided by ERIM for each sample. This will normally be range-dependent. The last column is used to identify which particular sensor is being used to measure the data at any particular point in the field of regard. The physical units for the above measurements are seconds, feet, feet, and degrees.

11

Filename: 11
Vehicle ID: 11
Status Code: 467
Encounter Time: 11.200000
Vehicle Width: 7
Vehicle Length: 23
Number of Estimates: 81

Time	X	Y	Yaw	Error	Sensor
0.00	-0.38	0.20	-1.62	3	0
0.10	6.77	-0.03	-3.27	3	1
0.20	14.06	-0.34	3.32	3	1
0.30	21.46	-0.24	-3.76	3	1
0.40	28.40	0.07	-4.46	3	1
0.50	35.58	-0.34	-2.82	3	1
0.60	42.47	0.01	-0.71	3	1
0.70	49.55	-0.08	-3.42	3	1
0.80	56.07	0.18	2.36	3	1
0.90	62.89	-0.20	1.40	3	1
1.00	69.77	-0.23	-0.98	3	1
1.10	76.58	-0.04	2.07	3	1
1.20	82.34	-0.03	-1.33	3	1
1.30	89.07	0.32	5.12	3	1

Figure 4-25. Text Display of Track File.

Graphical Plotter

The plotter element allows users to graph track file data or other calculated results (e.g., Kalman filter outputs or smoothed track file data) as a function of time or to cross-plot any two variables. Overlays of multiple data sets on one plot can also be obtained. The plots can be saved to disk or printed. An example plot window is seen in Figure 4-26.

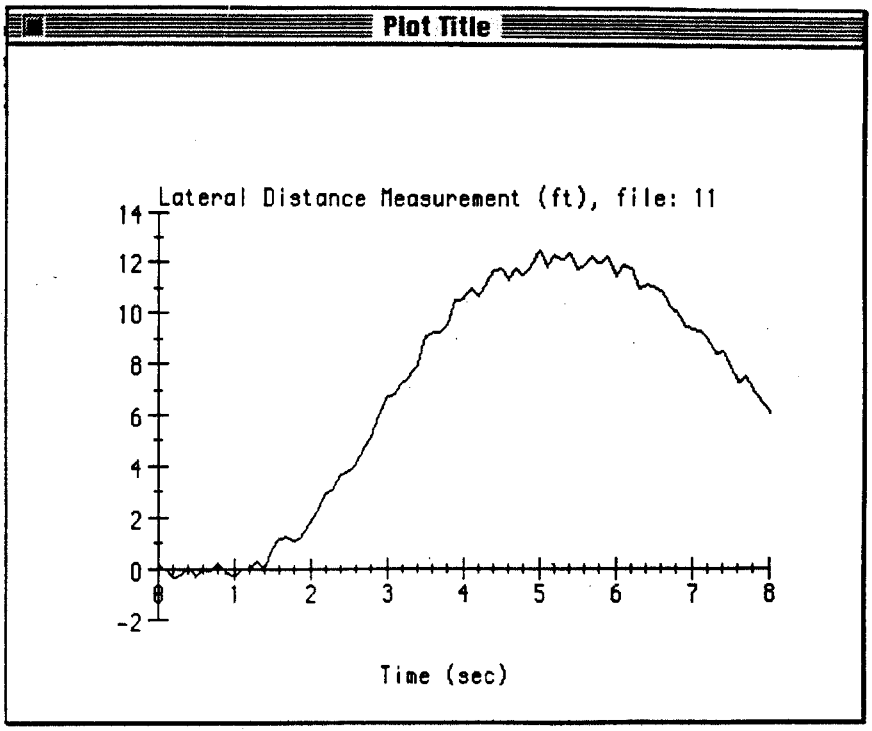

Figure 4-26. Example Plot Window.

Animator

The animator module permits the user to animate sequences of track files on-screen as a simplified overhead view of traffic flow. This feature can be useful for viewing the basic dynamic interactions of vehicles as they move through the field of regard. The on-screen movie can be started form various reference times and will play continuously until interrupted by the user. A time-base or clock reference is also seen on screen during the animation to help identify and locate which track files of the database are currently being observed. Figure 4-27 shows a portion of the animator on-screen window display.

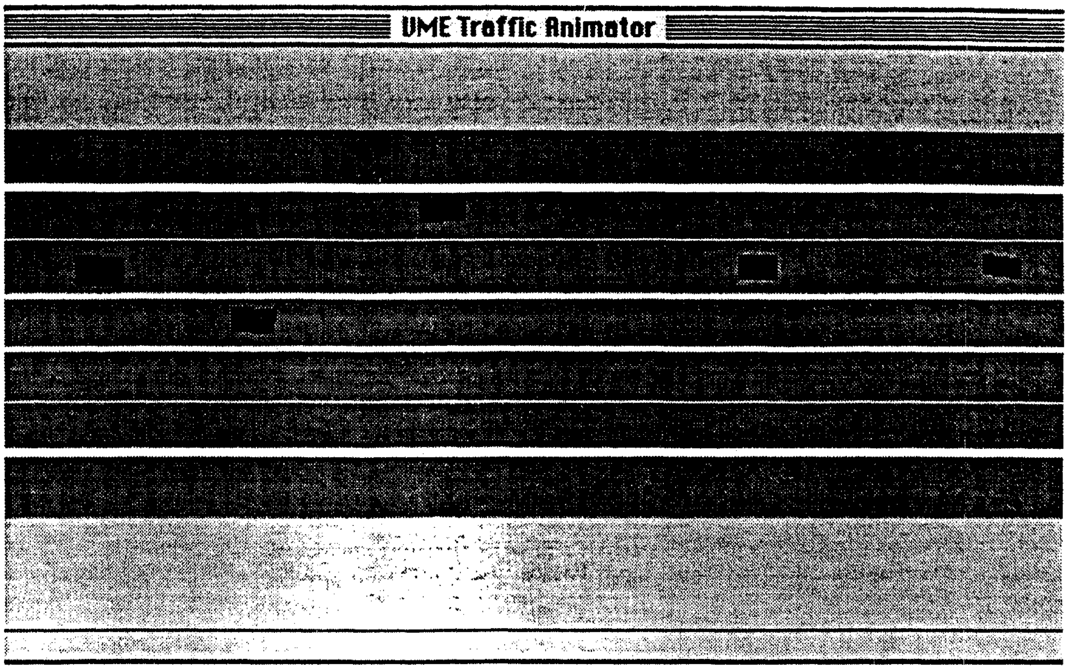

Figure 4-27. On-Screen Animation Display.

Kalman Filter

This module applies a Kalman filter calculation to the selected file and extracts five additional driver/vehicle response time histories. The Kalman filter calculation is tuned to vehicle size and speed using information contained in the track file header record (vehicle length parameter and initial speed of entry into the field of regard). The Kalman filter calculations utilize a simple three degree-of-freedom vehicle model composed of lateral translation, longitudinal translation, and yaw (heading) rotation. In addition to the three signals being directly measured by the sensor package and constituting the primary track file information [forward displacement (x), lateral displacement (y), and heading angle (psi)], the Kalman filter also estimates five additional driver/vehicle system response variables. These are: forward speed, lateral speed, yaw rate, front wheel steer angle, and longitudinal acceleration. The latter two response variables, front wheel steer angle and longitudinal acceleration, represent driver control response inputs to the vehicle required to achieve the gross vehicle motions reflected in the x, y, and heading measurements. The output from the Kalman filter also produces improved estimates for the three measured states (forward position, lateral position, and heading angle. Consequently, the total output from the Kalman filter calculation is an eight-state vector comprised of: longitudinal vehicle position, lateral vehicle position, vehicle heading angle, forward vehicle speed, lateral vehicle speed, vehicle yaw rate, front wheel steer angle (driver), and longitudinal acceleration (driver). An example plot comparing heading angle versus its Kalman filter estimate is shown in Figure 4-28.

The Kalman filtering concept enhances the conventional measurement process by utilizing knowledge of the basic dynamics of the system being measured. In the VME application, the approximate dynamics of cars, trucks, or other ground vehicles are used by the Kalman filtering calculation In practice, this results in a type of compromise between direct measurements of the system and corresponding state predictions provided by the dynamic model of the measured system. If a particular state measurement process is particularly noisy and differs significantly from state estimates predicted by the dynamic model, the noisy measurements are given less weight as part of this continuous measurement - model compromise. On, the other hand, if a measurement process is known to be particularly accurate, the measured values are given far more weight than the model predictions and would then dominate the output provided from the Kalman filter calculation.

Defining an appropriate dynamic model, identifying expected errors in the sensor measurements, and assigning accuracy levels to the dynamic model being used, is generally referred to as "tuning" the Kalman filter. For the VME program, a certain amount of tuning has been done and is built into the default calculation. However, future use of this feature will permit users to further tune and refine the filtering calculation parameters away from the default settings. Also, users can always export the raw track file data to disk and then import it into their own Kalman filtering program such as MATLAB or an equivalent analysis program for more advanced applications.

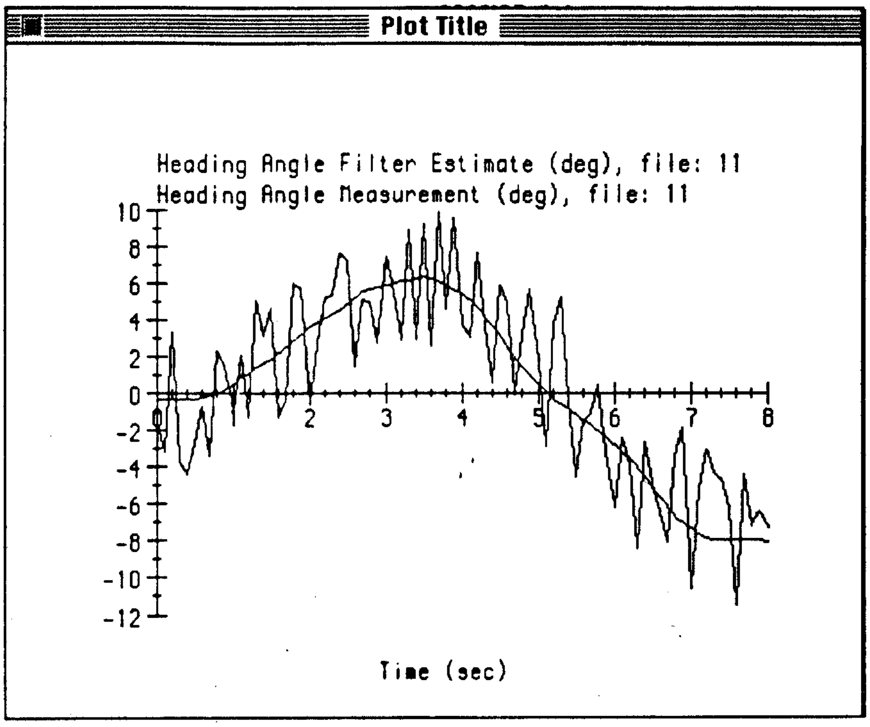

Figure 4-28. Plot of Kalman Filter Calculation vs. Raw Measurement.

Special Calculations

Specialized calculations applicable to larger groups of track files are grouped under this program feature. At the present, two special calculation modules are available — one for detecting crash events, the other for computing large quantities of inter-vehicular spacing information (time histories of range and angle-of-attack between all vehicles within the field of regard at each sample time).

Crash Detection

This special calculation module calculates the inter-vehicular gap between adjacent vehicles at each instant of time within the field of regard. As the traffic flow proceeds, any gap calculation that falls below a specified threshold set by the user will be tagged and recorded to disk file for later review. An option to view an animation of the traffic as the crash detection calculation occurs is also available. Under this option, a crash event will cause the two intersecting vehicles to change color and thereby assist the user in detecting the computed crash event on screen. See Figure 4-29.

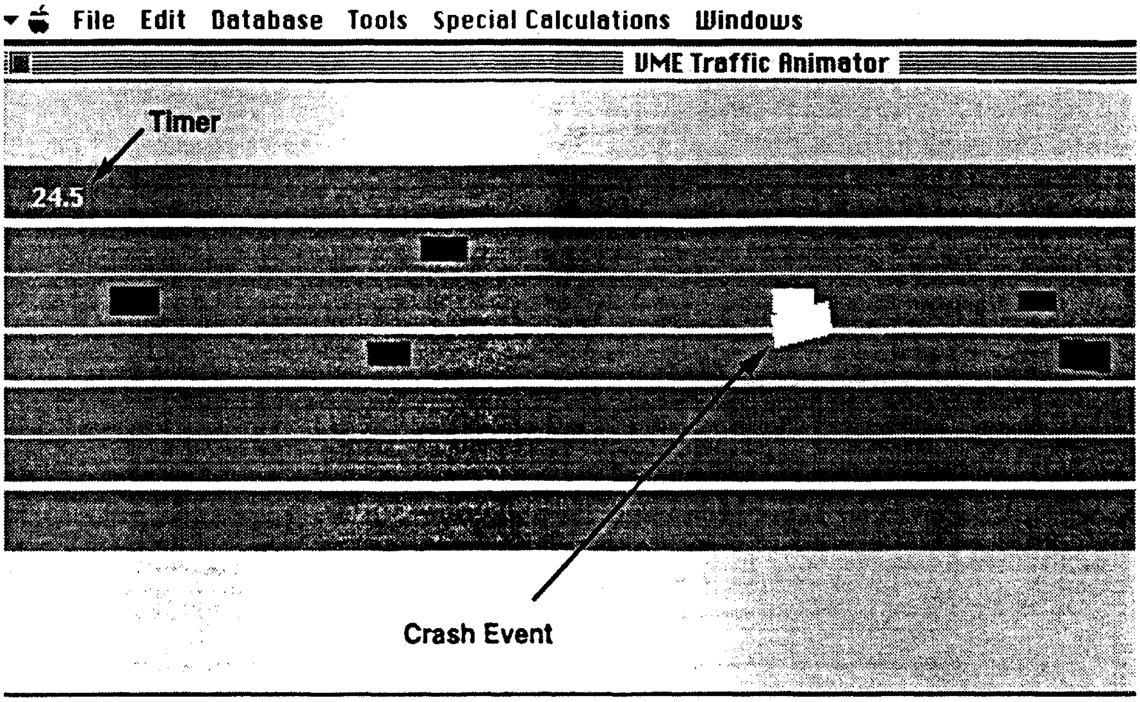

Figure 4-29. Crash Detection Calculation and On-Screen Display.

Computation of Inter-vehicular Variables

Under this module, the range and angle-of-attack between any two vehicles within the field of regard are computed and exported to a disk file. This computation occurs for each vehicle pair and for each sampling time during a time period specified by the program user. For example, in Figure 4-30 the range and angle-of-attack variables are seen depicted for one vehicle relative to all other vehicles in the scene. These variables, plus the corresponding variables calculated for every other vehicle in the scene, are exported to disk file at each sample time under this batch mode computation.

These exported data can then be further processed with statistical analysis software or other programs to create histograms of different variables. It can also be used to provide a digital record of the traffic environment — for example, as a numerical characterization of traffic flow for studying the response of a simulated crash avoidance package under various traffic conditions.

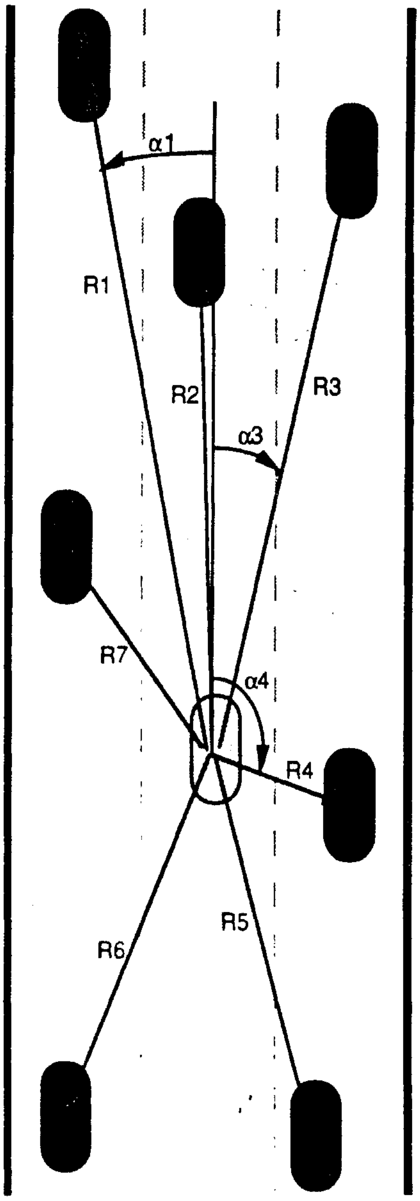

Figure 4-30. Range and Angle-of-Attack Variables Calculated as Inter-Vehicular Variables.

4.2.4 Example Usage

An example usage of the VME-DS software is described below to help communicate the basic nature of the program as it currently exists. The example shows how to open a database, and select a group of files for viewing with the text editor and the animator tool, and then activate the Kalman filter calculation and plot several of the resulting data.

— Opening a database and searching —

The basic software interface to the program user is a typical Macintosh or Windows-like environment with pull-down menus, dialog boxes, scrollable lists, etc. The basic intent is to provide a simple "point and click" environment for performing most operations.

Upon launching the program, the user first needs to select and open an existing database to search/analyze. Figure 4-31 shows the "**Open Database ...**" item under the "**File**" menu. By choosing this item the user is presented with a standard Mac dialog box containing a list of available databases. After selecting a database in this list, it is opened and becomes available for searching/processing.

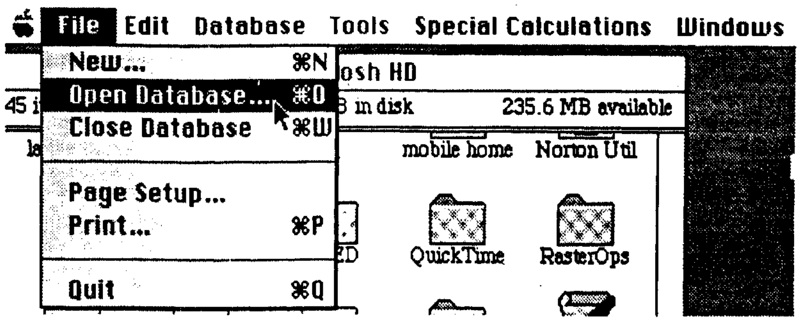

Figure 4-31. Opening a Track File Database.

This selection will then produce the user dialog box seen in Figure 4-32 that allows the user to search the current database by time or other information available in the vehicle header record (e.g., vehicle length, width, ID, etc.).

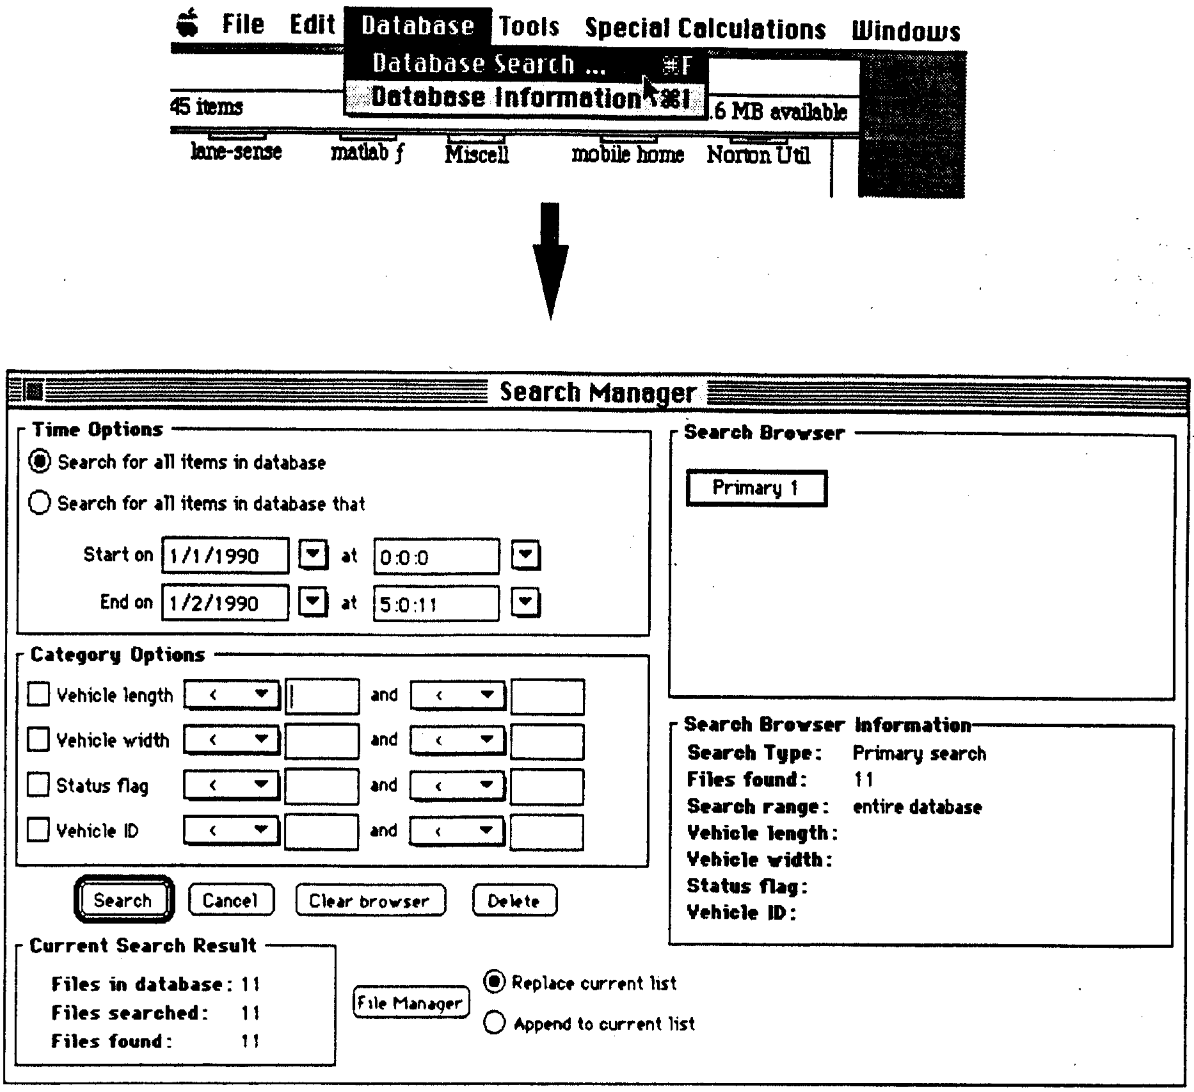

Figure 4-32 Searching the Current Database.

The dialog items appearing in Figure 4-32 permit the user to set the starting and ending date/time information for the search and to set basic Boolean operations for use in the search. This search dialog may be modified in subsequent versions to include vehicle header information fields that may be added in future versions.

By clicking on the **Search** button in the dialog, the user is presented with the following dialog window seen in Figure 4-33 showing the results of the search as a scrollable list of Vehicle ID names in the top portion of the dialog. This search results list may be further modified by the user, or copied in its entirety, to a second list appearing in the lower portion of the window. This second lower list represents the final list of files that the user has to work with during this particular search session. Individual files, or all files, from this final list may the be selected for viewing, graphing, or animation. The button functions at the bottom of the file manager dialog window directly affect those files selected in the second lower scrollable list.

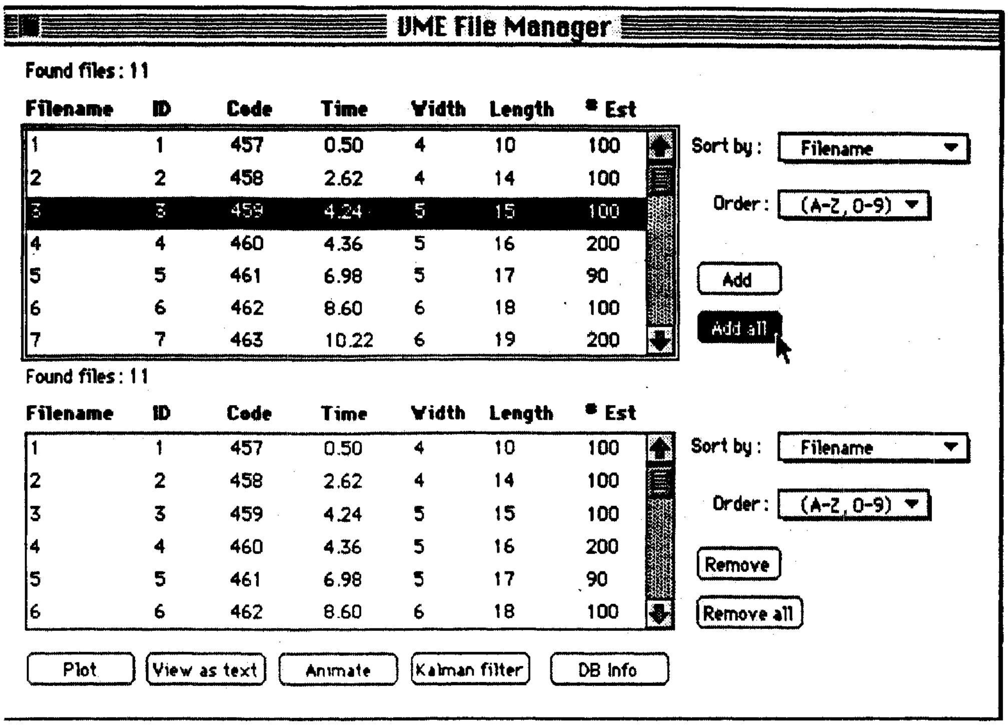

Figure 4-33. Refining the File List for Processing.

— *Viewing, Graphing, and Animating Selected Data* —

For example, to now see the contents of file 11, the user clicks on file 11 in the lower scrollable list to select it. To view that file as a text file the user would then click the **View as text** button at the bottom of the window. (See previous Figure 4-25 as an example result.) To graph certain variables from the track file, the **Plot** button would be used. (See previous Figure 4-26 for an example plot.)

To animate all the files in the lower list, the user clicks on the **Animate** button. Figure 4-34 shows an example of the resulting on-screen animation . If animation of only two vehicles were desired, the user would clear the lower list and then move only those two files from the top list to the lower list prior to clicking on the **Animate** button. The playback speed can range from slow motion to faster than real time, depending on the playback speed selected by the user.

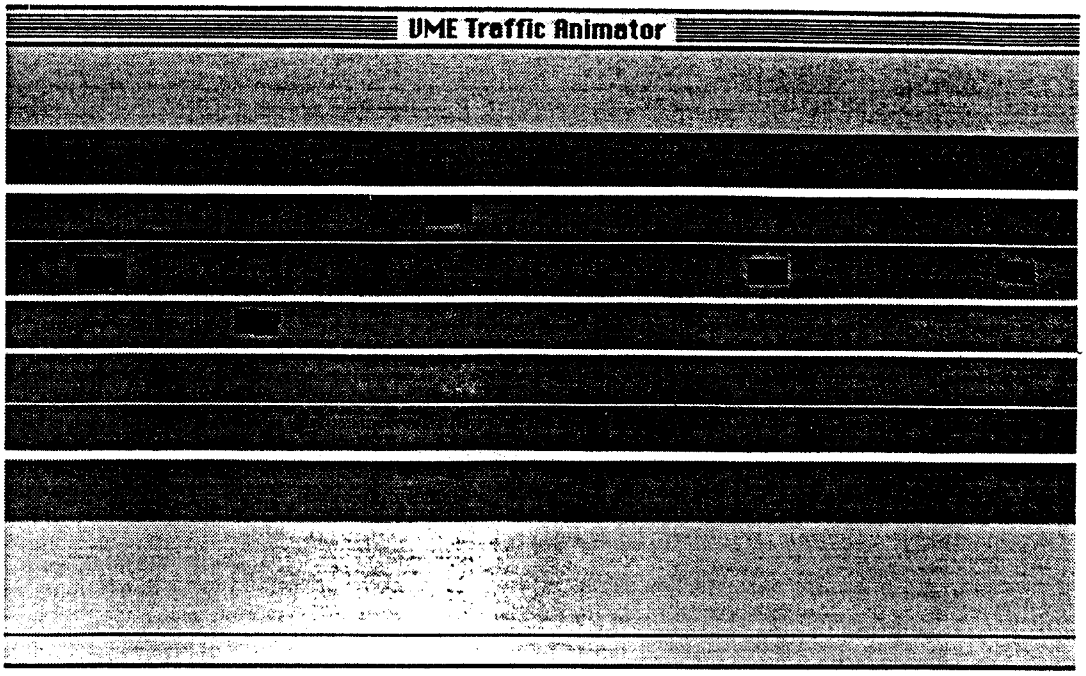

Figure 4-34. Snap-Shot of Animated Vehicle Flow.

— Kalman Filtering of Selected Data—

The **Kalman Filter** button in the file manager dialog window permits the user to select one or more files for Kalman filter processing. When this button is clicked, a secondary dialog window now appears (see Figure 4-35), which presents the user with a list of variables for viewing. This list corresponds to the original set of measured data (x, y, heading angle) PLUS the augmented system responses produced by the Kalman filter calculation. These augmented variables provided by the Kalman filter include estimates for the original measurements (x, y, heading angle estimates) PLUS estimates for forward speed, lateral speed, yaw rate, front wheel steer angle, and longitudinal acceleration. Each of these eleven variables (3 raw measurements + 8 Kalman filter outputs) may then be cross-plotted or viewed in the same manners as described above.

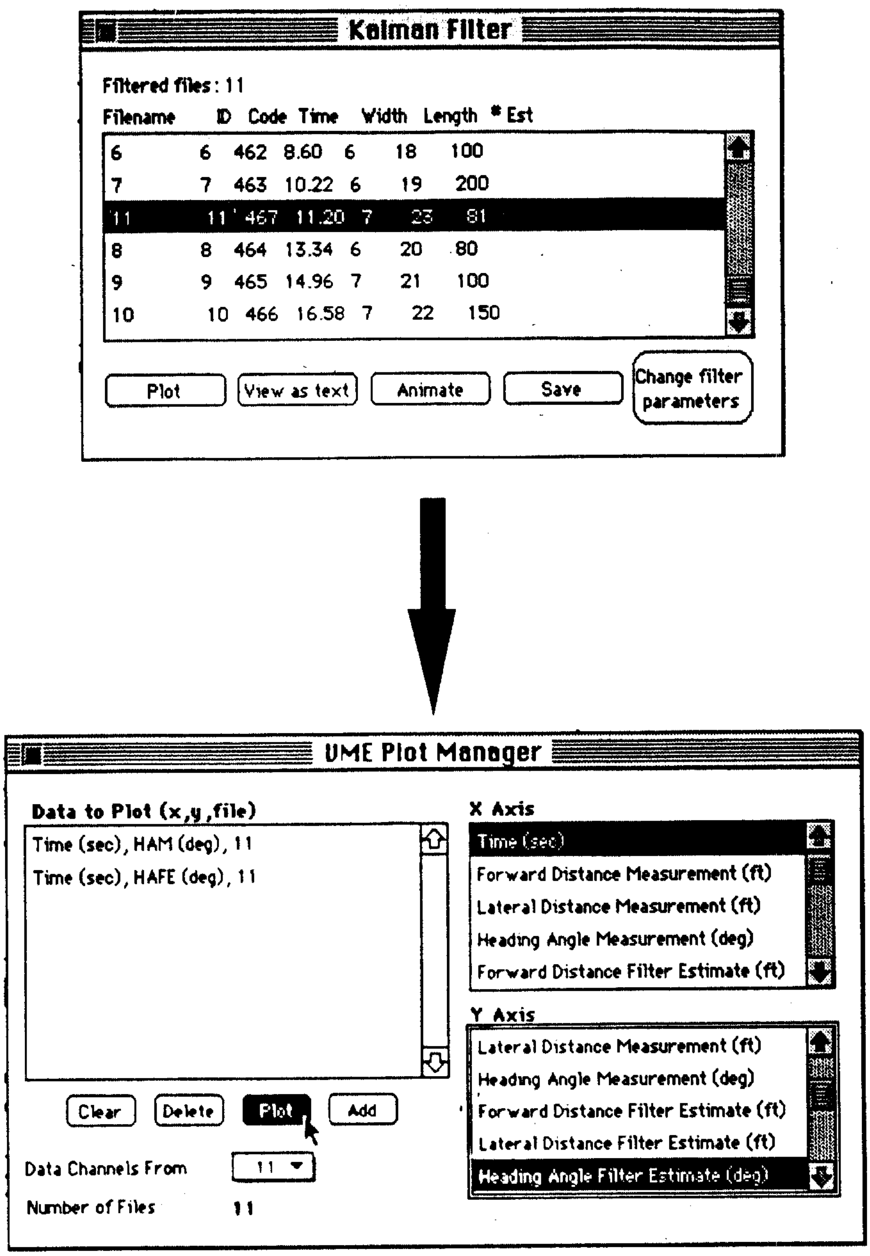

Figure 4-35. Kalman Filter Dialog and Plotting of Measured vs. Estimated Data.

By way of example, suppose that file 11 was selected for Kalman filtering and the user wished to compare the raw measurement data contained in the original track file with those same system responses estimated by the Kalman filter. The user would then select file 11 in the file manager list and click the Kalman filter button. The user could then select from the list of variables in the Kalman filter dialog. Figure 4-36 shows two example graphs from file 11 comparing the measurements of lateral displacement (y) with its Kalman filter estimate (top graph) and the measurement of heading angle with its Kalman filter estimate (lower graph)

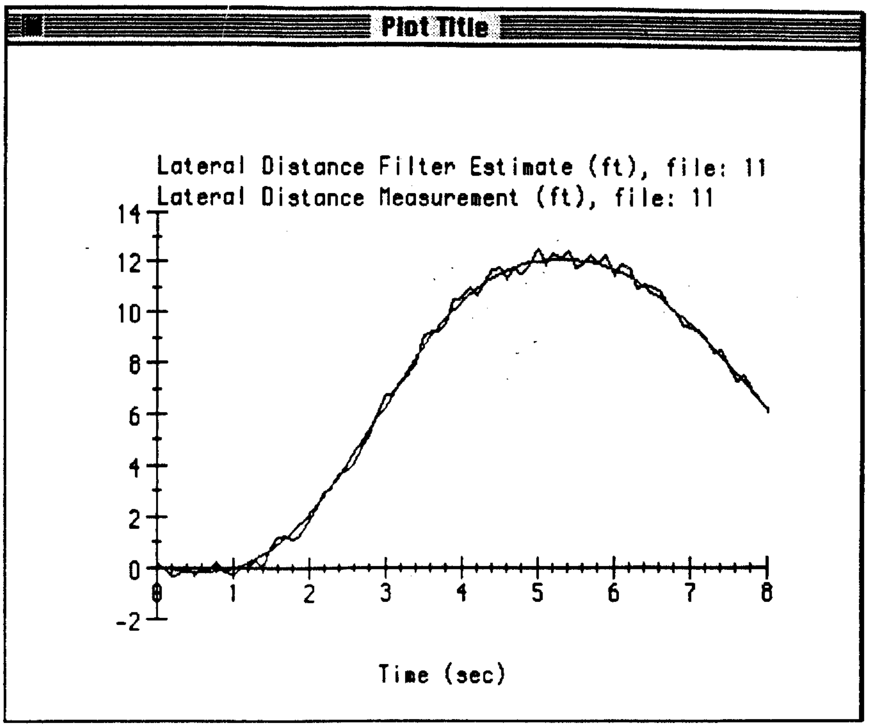

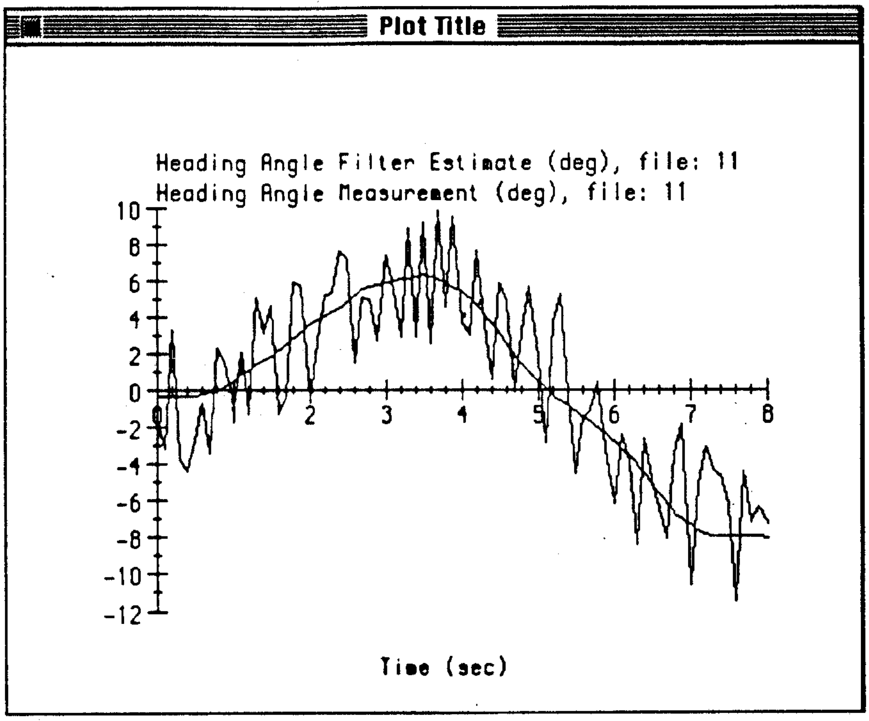

Figure 4-36. Plots of Measured vs. Estimated (Kalman Filter) Track File Data.

The user may also wish to examine estimates provided by the Kalman filter of control actions provided by the driver of that vehicle. The driver steering response is reflected in the estimate of front wheel steer angle. (See Figure 4-37.) The braking/throttle control response of the driver is reflected in the longitudinal acceleration estimate provide by the filter calculation seen in Figure 4-38. (This particular track file contains data for a vehicle performing a braking maneuver during a lane-change.) Consequently, information well beyond that provided by direct sensor measurements of vehicle location and orientation (raw track file data) are possible if Kalman filtering techniques can be correctly applied.

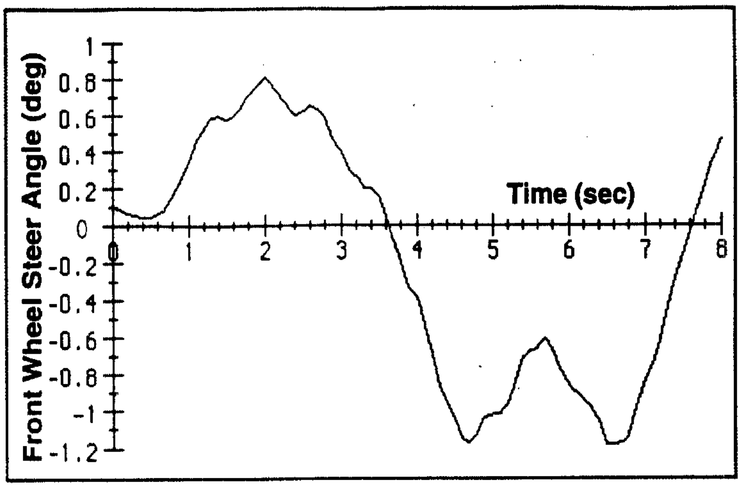

Figure 4-37. Kalman Filter Estimate of Driver Steer Angle Response.

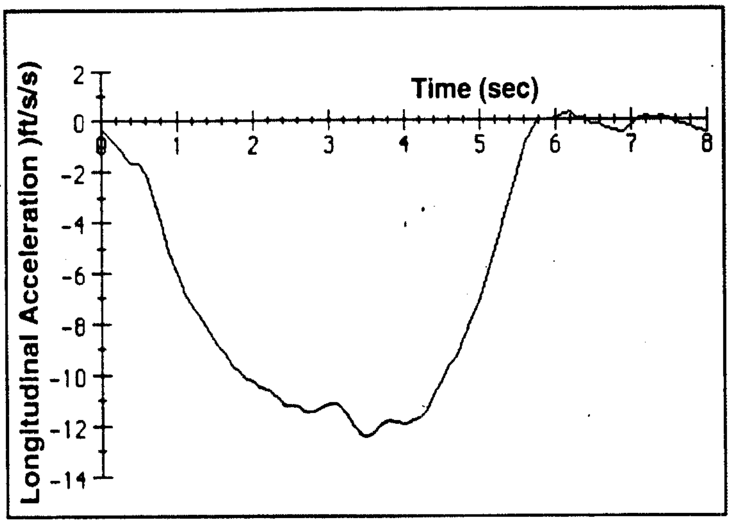

Figure 4-38. Kalman Filter Estimate of Driver Longitudinal Acceleration (Braking) Response.

4.2.5 Example Applications

The VME-DS software can be used to study a variety of traffic safety issues. These can include such topics as evaluating countermeasure technologies for crash avoidance, quantitative descriptions of traffic flow for testing ITS technologies, or human factors studies requiring quantitative information on headway spacings or lane deviations exhibited by drivers. Human factors specialists might also utilize information about frequency distributions of driver steering and braking control behavior under various traffic conditions. Using the VME-DS program to export these various data, a variety of research analyses can be conducted. Some examples are:

- *Crash Avoidance Research on the Potential Impact of Countermeasure Technologies*

This research area is concerned with examining specific crash incidents or near-misses recorded by the VME system and utilizing post hoc "what if" analyses to illustrate the potential of various technological assists in helping mitigate the likelihood of such incidents. For example, a recorded crash or near-miss incident could be re-played, using computer simulation of specific crash avoidance technologies and the same recorded VME traffic environment, to query whether or not a candidate on-board crash avoidance system could have helped avoid, or lessened the severity of, the actual crash or near-miss incident. Such analyses, spread across a variety of different recorded incidents, could then be used to evaluate the potential effectiveness of such systems and their likelihood for improving safety.

- *Quantitative Traffic Conflict Data for use in Safety "Bridging" Analyses*

Another potential research activity is the use of VME data to catalog and record information at specific traffic sites for assisting safety analyses that hope to someday "bridge" between observations of routine traffic behavior and the likelihood of crash events. For example, "time-to-collision" numerics are frequently used to estimate the likelihood or potential for increased crash events. However, time-to-collision data, or similar traffic-conflict numerics, are generally not available and can be very time consuming to obtain — usually through conventional video recordings and manual processing methods. The availability and use of VME data in this capacity could help to support and accelerate various bridging analyses now underway.

- *Quantitative Description of Traffic Environment for Testing/Evaluation of ITS Technologies*

Manufacturers of ITS equipment intended for use on-board vehicles as warning and/or control intervention devices would presumably have needs to test and evaluate their devices in realistic driving environments. If databases of regular traffic data were available from a VME system, manufacturers could then expose their systems to "electronic" traffic from the VME database(s) and evaluate how such system would likely respond under different traffic conditions. Costs and times associated with adjustment and tuning of such systems would be facilitated with the number of on-highway field trials being significantly reduced.

Various statistics or traffic environments could also be summarized in many cases by frequency distributions of traffic flow. For example, Figures 4-39 and 4-40 show simple histograms of angle-of-attack and side-range variables obtained from one-minute of a hypothetical (computer-generated) VME database traffic flow. This data corresponds to three lanes of freeway traffic all flowing in the same direction. Figure 4-39 shows simply the number of targets detected forward of each vehicle in a +/- 10 degree viewing angle (similar to a forward-looking radar).

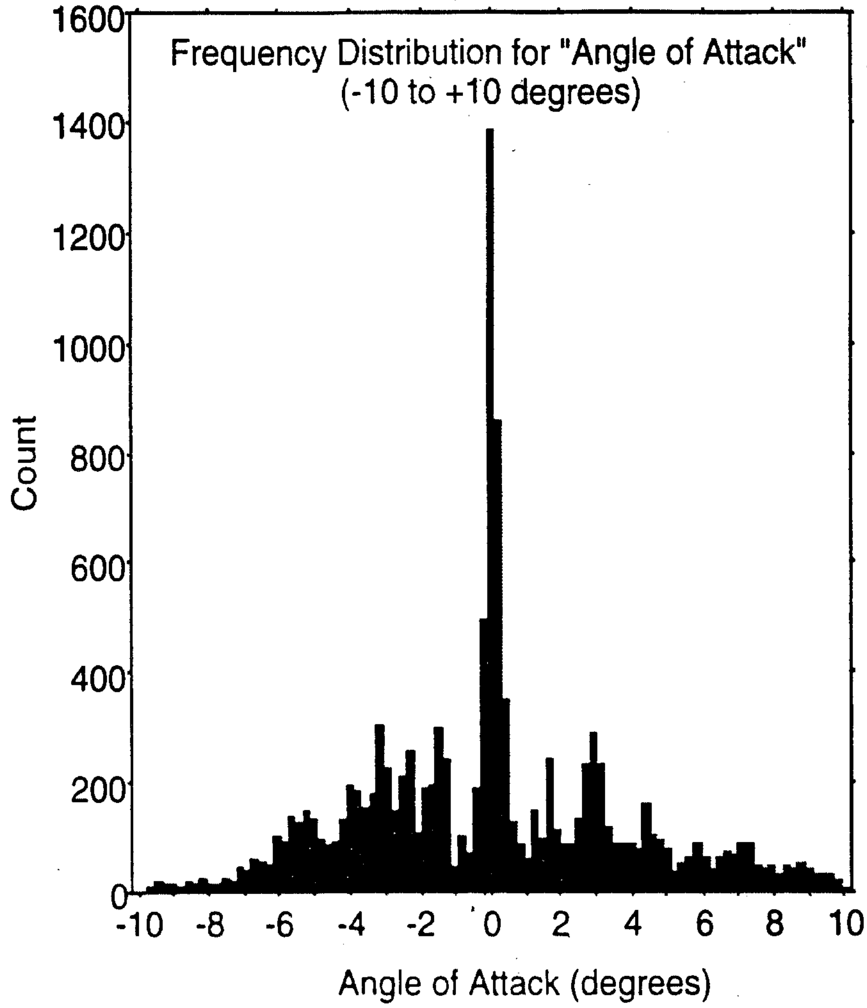

Figure 4-39. Histogram of target detections forward of each vehicle (1-minute of computer-generated traffic flow).

The count at 0-degrees angle-of-attack represents vehicles ahead of one another. Side-lobes of increased counts at plus and minus 3 degrees represent vehicles in adjacent lanes, but within the confines of the +/- 10 degree overall viewing angle. Manufacturers of crash-avoidance sensors could utilize this or similar types of information to initially gauge the expected operating environment for their sensor.

Similarly, Figure 4-40 shows a count or histogram of vehicles detected laterally (70 - 110 degree field of view, left and right sides) from each vehicle as a function of range. Since the lanes are each 12 feet wide, clusters appear at 12 and 24-feet intervals. The smaller count cluster in the vicinity of 16 to 18 feet of lateral range represents vehicles occasionally performing lane-change maneuvers during this one-minute of traffic flow.

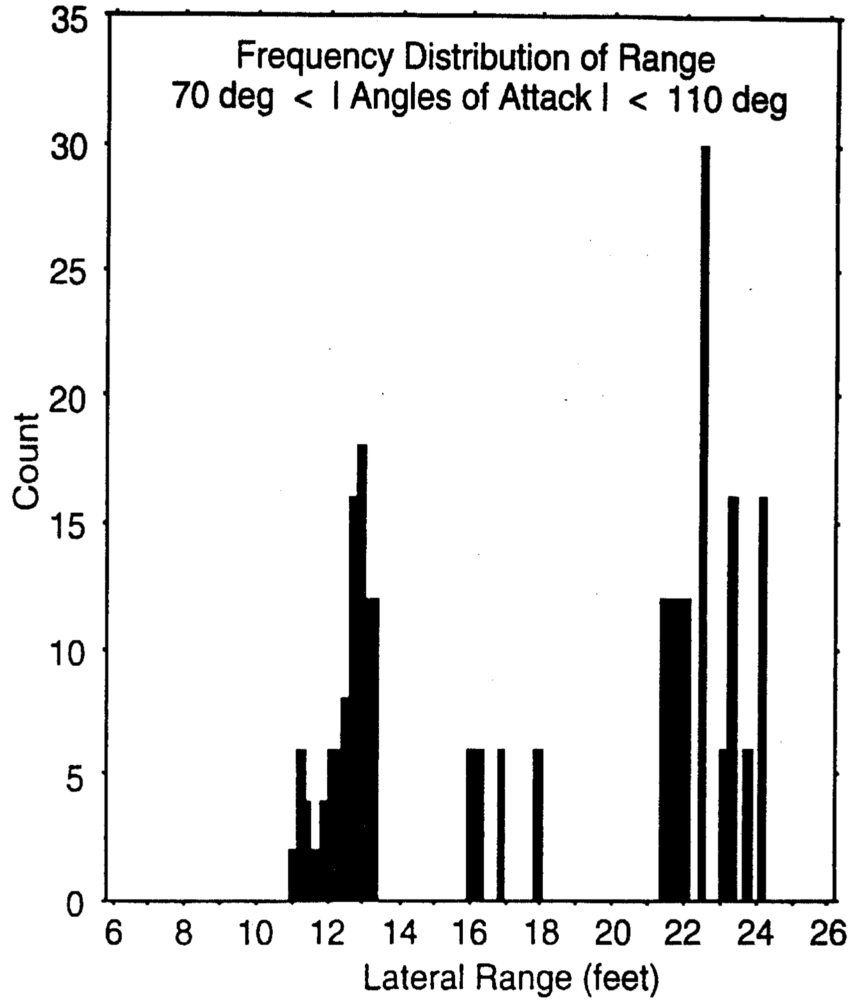

Figure 4-40. Side-Viewing Experience for 1-minute of Traffic Flow. (computer-generated traffic data)

- *Human Factors Studies on Driver Behavior*

Lastly, the area of human factors research and driver behavior can also derive useful information from a VME database and its associated software. As seen in previous sections, the Kalman filter feature of the VME-DS software permits the program user to extract additional driver-vehicle responses from the original raw track file data. Two of these are essentially driver control responses: front wheel steering angle and longitudinal acceleration (braking/acceleration). These control inputs represent estimates of how a driver must steer and brake/accelerate so as to achieve the trajectory (X, Y, and heading-angle measurements) recorded in the raw track file by the sensor hardware. These estimates of driver control behavior can then be exported and processed to obtain various numerics such as RMS measures of driver steering behavior or braking/acceleration behavior.

Similarly, lateral path deviations within a lane, or headway distances within different traffic environments, can also be exported and analyzed in a similar fashion. Average RMS levels of lateral "wandering" within a lane may be of interest, or comparable measures of driver steering activity. Likewise, headway distances and/or speeds used by drivers under different traffic conditions may also be of interest. By exporting such information from the VME-DS and conducting follow-up analyses with more specialized software packages, a wide range of analysis alternatives become available while also providing flexibility for continued availability and distribution of VME data.

5.0 RESULTS OF FIELD DEPLOYMENT AND TESTING

The VME-MS underwent a series of tests prior to and during deployment. Before deploying the VME-MS on ERIM's property facing a heavily traveled public roadway, however, a number of authorizations had to be obtained based on a layout for this specific site. The results of these activities are presented and discussed in the following five subsections.

Highbay Testing - The VME-MS was partially deployed in a four-story, highbay area that permitted initial functional and performance tests to be conducted in a very controlled setting. The results of these tests indicated that the system was functionally ready for deployment and that the short-range performance data was consistent with the theoretical projections.

Site Layout - The VME-MS was deployed on ERIM property along Plymouth Road in Ann Arbor, Michigan. This is a very heavily traveled section of roadway and in close proximity to a major intersection. The details of the field site deployment with a site drawing are briefly discussed.

Authorizations - The deployment of a laser-based device, on a 100 foot tower in a public setting required safety certifications and approvals from a number of public agencies. The processes for obtaining these approvals, and the resulting approvals, are summarized to demonstrate that the VME-MS was deployed in a safe and appropriate manner.

Sensor Station Installation - The process of physically deploying a VME-MS Sensor Station is reviewed and summarized. An Operations Manual providing greater detail regarding the installation of the Sensor Stations is provided in Appendix E.

Field Testing - Mechanical, electrical and functional tests were performed as part of the field deployment demonstration. The mechanical tests and associated analyses were performed and verified that the VME-MS was safe to deploy. The electrical and functional tests demonstrated that a single Sensor Station could acquire and partially process 3D-laser sensor data in real time. These tests were not completed because it became clear during the tests that the 3D-laser data quality was inadequate to justify any further testing.

5.1 Highbay Testing

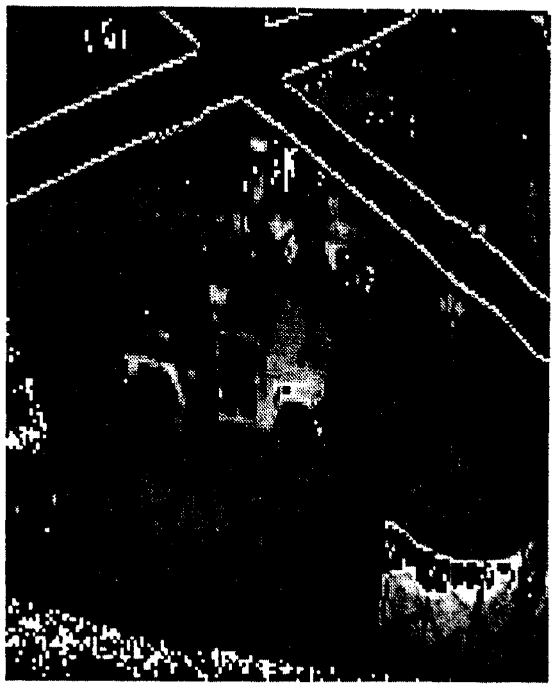

Figure 5-1. Reflectance Image

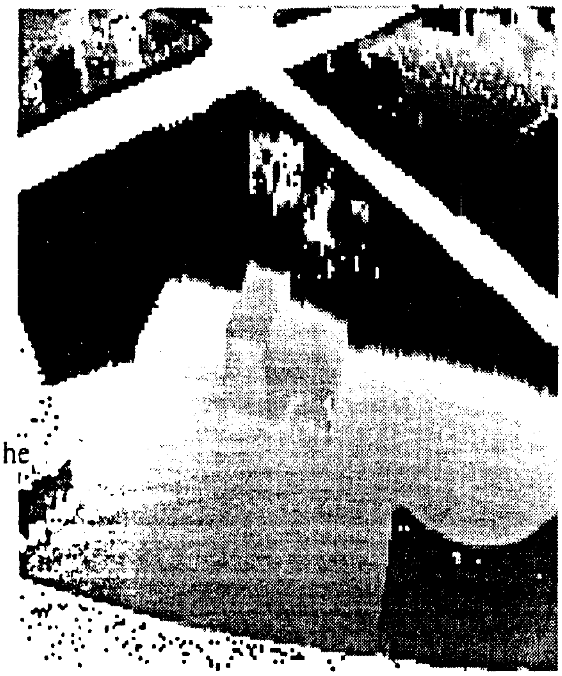

Figure 5-2. Range Image

Setting up a Sensor Station in ERIM's "highbay" environment allowed the sensor's performance to be analyzed in a controlled setting. The sensor could be raised to a height of 32 feet above the floor. Vehicles and objects could be placed at known locations in the sensor's field-of-view and components of our SCPU image processing could be tested. Figure 5-1 and Figure 5-2 show examples of an intensity image and a range image, respectively, collected in the highbay. Both figures show the roof support beams crossing the upper portion of the image and a jeep in the center of the field of view. Dark values in the reflectance image mean that the laser is reflecting off of a dark object or off of an object far from the sensor. The support beams are very close to the sensor causing them to reflect a lot of the energy back to the sensor. The range image gives modular range. The value 0 is coded as black while the value 50 feet is coded as white. There is a transition boundary in Figure 5-2 with white below and black above. This is where the modular range makes a transition from 50 to 0 feet. This transition zone is called the ambiguity boundary.

The only disadvantage of using the highbay was that the maximum range was nowhere near the range that the system would have to operate at when taken to a roadway site for the demonstration test On the other hand, we could perform several experiments before weather conditions allowed the tower to be erected outside. Several tests were conducted.

(1) Range accuracy was estimated. Range accuracy as determined in the highbay was consistent with the corresponding measurements obtained at Perceptron and predicted results for the sensors with reduced laser power.

(2) The tower motion detection procedure was evaluated. We used the intensity channel to track tower motion. Reflectors placed at known locations in the scene allow any tower motion to be estimated. The detectability of various types and sizes of reflectors was measured. Retro-reflecting tape could be seen in the intensity image, but the best results were obtained using the type of reflecting unit designed to be embedded in the road. This experiment would have to be continued when data is collected at full range. Figure 5-3 shows an intensity image containing retro-reflecting tape ("tape"), and a lane-marking embedded reflector.

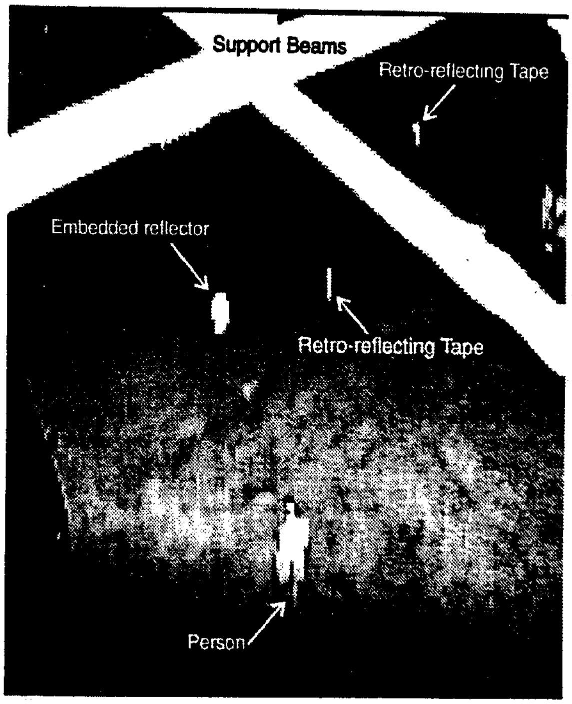

Figure 5-3. Reflectance Image with Retro-reflecting tape

(3) The range reference updating procedure was tested using real data. Test results showed some of the expected pattern range values increasing to 50 feet as you move along the floor and then the values drop off to zero. The transition boundary, however, showed unexpected characteristics. The 3D sensor technology can provide range modulo 50 feet. That is, the sensor can provide a range value whose error is an inch or so plus an integer multiple of 50 feet. ERIM has used models to predict the integer multiples of 50 feet for adding to the 3D derived range producing the "range-recovered image." When these models were applied to ERIM's experimental 3D sensors, the resulting error was an inch or so. Perceptron's implementation of the 3D sensor provides a random like value near ambiguity boundaries, which is independent of range. This occurs because their 3D sensor averages over the ambiguity boundary.

This effect can be seen in Figure 5-2 as a blurry transition from values near 50 feet to values near zero. Figure 5-4 shows the resulting error when our range model is used to predict the best integer multiple of 50 feet to add to the signal measured from Perceptron's 3D sensor. The near zero errors of most of the floor are displayed as a mid-level gray value. Negative errors are darker grays to black, while positive errors are lighter grays to white. There are several places where the error is significantly different from zero. The person in the lower center portion of the image is the most obvious. Because the person was walking through the scene, he was not part of the background model, thus causing him to show up as an error (with respect to the model.) This is the principle behind our vehicle detection algorithm, the road without vehicles is modeled. When vehicles produce data which is inconsistent with the model we reject the underlying hypotheses that there are no vehicles. Another region of the image where the range error is significantly different from zero is the ambiguity transition region, which is the arc across the image. The basic random nature of the 3D range information near the ambiguity boundary cannot be removed with a model. It becomes an apparent range discontinuity. Any algorithm that uses Perceptron's range channel over transition regions will have to address this artifact. This ambiguity problem is discussed further in Sections 5.5.2 and 6.3.

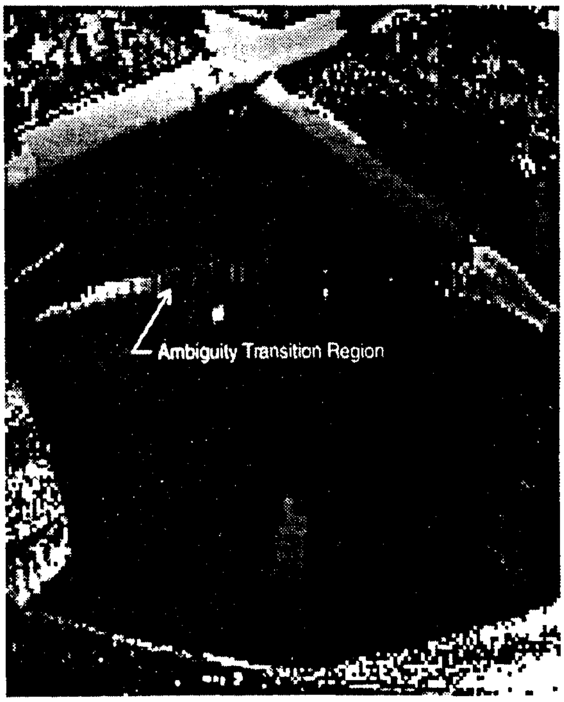

Figure 5-4. Range Difference Image

(4) The Sensor Station Operating System software was tested and verified to be functioning properly.

5.2 Site Layout

The layout of the Plymouth Road, Ann Arbor, Michigan site is shown in Figure 5-5. In general, a typical VME-MS site should support the ingress of a typical 4000-pound tow vehicle and its 2300-pound trailer/towers. Soil conditions should enable the installation of six (two per guy point) auger-style earth anchors per sensor station. A set-back distance of 32.5 feet due to the minimum guy radius at the ground plus any additional highway easement set-back distance is required. Hilly terrain can be accommodated, but it puts additional demands on towing, trailer leveling, and guy wire length fabrication.

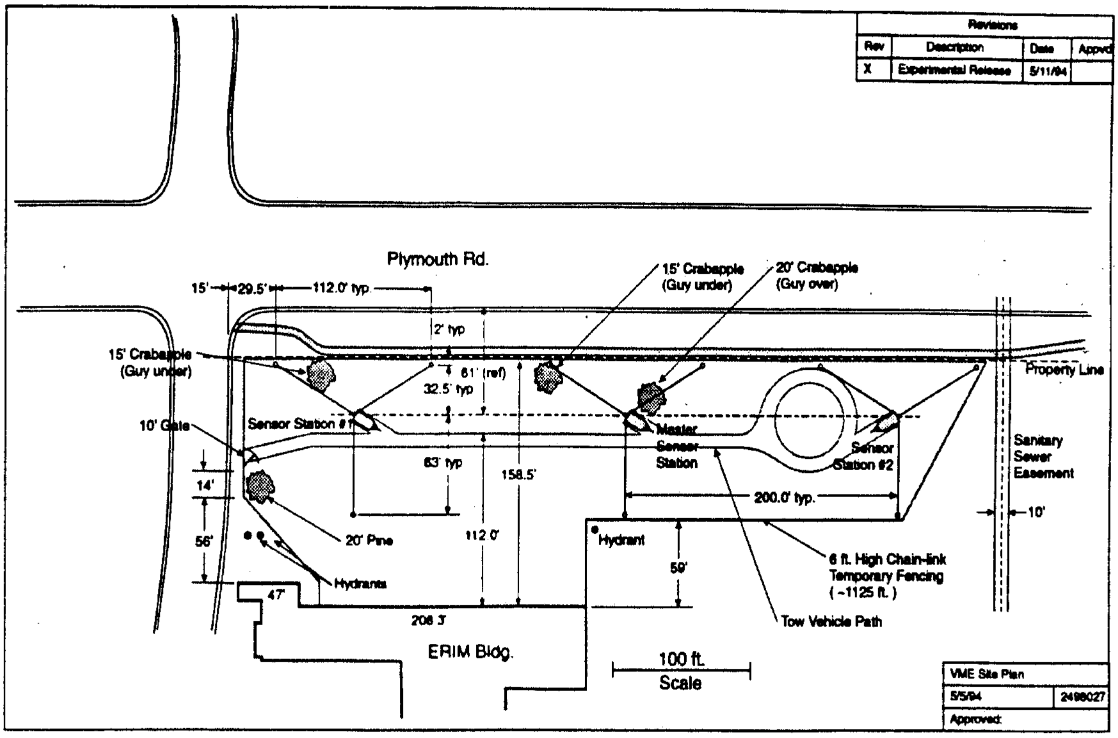

Figure 5-5. VME-MS Site Layout

A number of measures were taken in Ann Arbor to ensure the public's safety and that of the equipment. Specifically a 6-foot high chain link fence was installed to establish a closed perimeter within which the Sensor Station Assemblies (SSAs) were deployed. In addition, video cameras and intrusion detectors were placed in the vicinity of each SSA. The video cameras were connected to ERIM's building surveillance system and continuously monitored by ERIM's security personnel. The intrusion detectors were connected to a warning light that was located at an ERIM security station that overlooked the test site. Several signs designed to inform the public about the purpose of the tests and to indicate that the site was under constant surveillance were affixed to the perimeter fencing.

5.3 Authorizations

The deployment of a laser-based device, on a 100-foot tower, in a public setting required safety certifications and approvals from a number of public agencies. In anticipation of deploying the VME-MS during the spring and summer of 1994, ERIM started the process of obtaining the necessary deployment approvals from the appropriate agencies in February 1994. All approvals for an activity such as this are site specific, and as described in the previous section, ERIM chose to deploy the VME-MS on its property facing Plymouth Road. The processes for obtaining these approvals, and the resulting approvals, are summarized in the following sections. The submissions, approvals and receipt acknowledgments are provided in Appendix D.

5.3.1 City of Ann Arbor, Michigan

On April 11, 1994, ERIM submitted a petition for minor modifications to the City of Ann Arbor, Michigan Planning Department. This petition was essentially a request for a building permit to temporarily install three, VME-MS Sensor Stations on ERIM's property facing Plymouth Road. Concurrently, but separately, ERIM also submitted a request to the Zoning Board of Appeals for setback and height limit variances. These petitions resulted in two public hearings.

On Tuesday night, May 17, 1994, representatives from ERIM and UMTRI presented their plans for deploying the VME-MS Sensor Stations to the Ann Arbor City Planning Commission. The Planning Commission unanimously voted to approve the Site Plan for Minor Modifications. The following day, the same representatives presented the deployment plans to the Zoning Board of Appeals, who also unanimously granted the setback and height variances. All concerns relative to the nature of the project, structural integrity, public notification, duration, and public safety were satisfactorily addressed. It should be noted that both city groups were pleased to see such an interesting project of national importance being conducted in Ann Arbor.

5.3.2 Federal Aviation Administration (FAA)

The height of the VME-MS towers, 100 feet, suggested that we discuss our deployment plans with the FAA. The FAA indicated that they would like to review out deployment plans, and in particular the geographic locations for our towers. On March 14, 1994, ERIM submitted a "Notice of Proposed Construction or Alteration", FAA Form 7460-1, to the FAA's Great Lakes Region in Des Plaines, Illinois. The FAA reviewed our plans to determine if the towers would be an obstruction to air navigation, if additional markings or lights would be required and if FAA required any further notice of ERIM's plans. FAA's determination, dated May 27, 1994, was that: 1) ERIM's towers would not impose any obstruction to air navigation, 2) no additional marking or lights would be required, and 3) the FAA did not require any further notices.

5.3.3 Center for Devices and Radiological Health (CDRH)

In accordance with the Radiation Control for Health and Safety Act of 1968 (Title 21, Code of Federal Regulations, Subchapter J), ERIM submitted a Model Change Report on the VME-MS to the CDRH on May 12, 1995. Each Sensor Station of the VME-MS includes a certified, by Perceptron, Class IIIb LASAR Datacamera™ which is installed and deployed in such a manner that the radiation accessible during operation is Class I.

Based on ERIM's analyses performed with the assistance of a laser-safety consultant, ERIM has certified in its Model Change Report that the VME-MS is eye safe under the following conditions.

1. To the naked eye beyond five (5) feet from the exit aperture of the 3D laser sensor under normal VME-MS operating conditions;

2. To someone viewing the VME-MS through 7X binoculars at a range of thirty three (33) feet from the aperture of the 3D laser sensor under normal VME-MS operating conditions, and

3. To the naked eye beyond 103 feet from the exit aperture when the VME-MS nodding mirror is not operating and the exposure is continuous for up to ten (10) seconds.

5.4 Sensor Station Installation

After selecting a site with the qualities described above, a number of site-preparation activities can be completed prior to delivering the sensor stations to the site. An initial layout of the site should be performed to verify adequate space for fencing, guy anchor points, the tow vehicle path, and clearances to permanent site fixtures. Tower and guy point locations can then be measured and marked and the earth anchors installed and tested.

Two auger-style earth anchors must be installed at each guy point location. Each anchor is screwed into the earth at the average guy wire angle (approximately 45 degrees) and then proof-tested to 2,500 pounds using a special hydraulic load testing apparatus. Copper clad grounding rods are then driven at each guy point and each tower location point.

After the trailers are delivered to their predetermined locations they must be leveled and anchored using duckbill-style earth anchors. After anchoring of the trailer is completed the LSH and its cables may be installed on the adjustable mount on the tower and the appropriate azimuth and elevations settings selected and locked down. Next, the tower outriggers are rotated into position and the tower may be deployed and stowed per the VME-MS Operators Manual (see Appendix E).

5.5 Field Testing

Mechanical, electrical and functional tests were performed as part of the field deployment demonstration. The mechanical tests and associated analyses were performed and verified that the VME-MS was safe to deploy. The electrical and functional tests demonstrated that a single Sensor Station could acquire and partially process 3D-laser sensor data in real time. The tests performed and their results are discussed in the following two sections.

5.5.1 Mechanical Verification and Testing

The mechanical design approach for the VME-MS maximized the use of commercial off-the-shelf equipment. Accordingly, thorough verification and testing was performed to insure the safety of the equipment operators and of the general public during the deployment of the VME-MS. Detailed analysis, design and modifications of the tower, trailer and associated mechanisms were performed by ERIM's in-house engineers and technicians to increase the load capacity, and hence the safety factors, of these components. Additionally, the services of an independent civil engineering firm were obtained to verify the conformance of the tower design with standard local and national building requirements. In the end, a compromise was reached that balanced ERIM's stringent self-imposed safety requirements, the commercial off-the-shelf design approach, and program schedule requirements.

This compromise required that the towers only be deployed when the wind speed at the deployment site was no greater than 25 mph. However, to provide protection against a sudden storm, the tower restraint system of guy wires and their associated earth anchors and tower attachment brackets were verified to withstand a wind speed of 75 mph. Also, because of the danger of additional loads due to accumulated ice, deployment was further restricted to temperatures above the freezing point.

For purposes of verification, the VME-MS was grouped into three main assemblies of components. These assemblies were the tower assembly, the support frame assembly and the guy wire assemblies. The tower assembly included the four, 30-foot long tower sections and their associated cables, eyebolts, pulleys and fasteners that comprised the deployment mechanism. The support frame assembly included the rectangular support frame that transferred the tower loads to the trailer, the frame-support outriggers and all associated fasteners. The guy wire assemblies included the guy wires, the welded guy wire attachment lugs at the top of each tower section, the earth anchors, and all associated clamps and fasteners. The analysis, design modifications, and testing of these three assemblies are described below.

Safety Factors

All components affecting personnel safety were approved as follows. Components that had a manufacturer's working load rating (generally 20 percent of the breaking load, or 5X safety factor) were qualified as safe by assuring that the rated load was not exceeded when the tower was exposed to a 25 mph wind. These "5X" components included the wire rope cable, pulleys, and eyebolts. Components without a manufacturer's rating were subjected to a proof test load equal to two times (2X safety factor) the load calculated for a wind speed of 25 mph. This 2X safety factor is equivalent to the condition when the VME-MS is subjected to a 35 mph wind. These 2X components included the welded "A-Frames" for attaching the pulleys to the tower structure and all welded joints on the support frame and guy wire assemblies.

Tower Assembly Verification

Tower assembly analysis was initiated by inspecting the tower structural design and the design of its cable-and-pulley deployment mechanism. Next, equations (Figure 5-6) that described the load conditions of these components, as a function of guy loads, payload weight, and tower section weight were prepared. Finally, a table (Figure 5-7) of all load bearing components and their respective loads and safety factors was prepared in a spreadsheet format. The results showed that the load capacity of the commercial off-the-shelf towers was adequate for static loads only and that no additional load-bearing capacity was available for wind-induced loads.

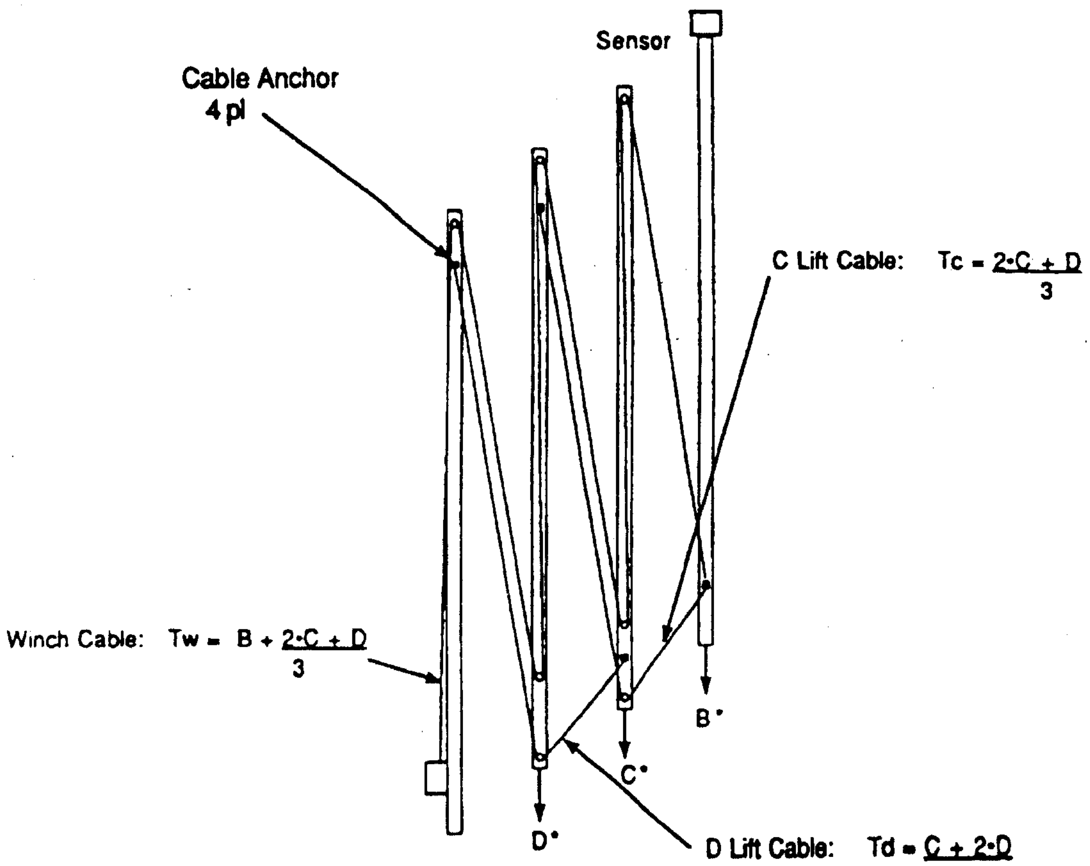

* Includes vertical guy loads, payload, and tower section weight

Figure 5-6. Tower Assembly Load Diagram.

	Load #	Limit #	Rated/Tested	Safety Factor
Main Cable				
Main Cable Tension	799	740	rated	0.93
Main Pulley Anchors	1597	3200	tested	2.00
Main Pulley	1597	1750	rated	1.10
Eye Bolt	1597	1500	rated	0.94
Lift Cables				
Section C				
Cable	352	800	tested	2.28
Cable Anchor	352	800	tested	2.28
Pulley Anchor	703	1600	tested	2.28
Pulley	703	1750	rated	2.49
Eye Bolt	703	800	rated	1.14
Section D				
Cable	340	700	tested	2.06
Cable Anchor	340	700	tested	2.06
Pulley Anchor	679	1400	tested	2.06
Pulley	679	1750	rated	2.58
Eye Bolt	679	800	rated	1.18
Tower Pivot Load	1456	3200	tested	2.20

Figure 5-7. Tower Assembly Load Matrix

The effect of wind-induced loads on the safety factors was determined by expressing the guy loads as a function of wind speed. The analysis showed that as the wind speed increased, the restraining tension in the guy wires induced significantly large downward loads on the tower's deployment mechanism. Because these wind loads reduced the component safety factors below acceptable values, design modifications were required. Trade-offs between maximum wind speed and tower modification complexity and cost were repeatedly iterated using the spreadsheet.

Modifications to the tower assembly included (a) installation of stiffening bars at the tubular tower rungs that had not been previously strengthened by the manufacturer, (b) installation of larger diameter pulleys having a greater load capacity, (c) installation of high-strength, hardened eye-bolts and (d) replacement of the secondary cables (lift cables) with larger diameter wire rope

Proof-test loads were applied to the tower assembly by restraining the tower sections one-to-another with high strength nylon rope and then extending the tower sections against these restraints. The integral electrical winch, used to extend the tower sections under normal circumstances, was used to generate the proof-test loads under the restrained conditions.

Support Frame Assembly Verification

The support frame assembly was analyzed for two separate loading conditions. The first condition occurs during tower "tilt-up" when the support frame assembly must provide the reaction forces in response to the loads applied by the manually operated "tilt-up" winch. The second condition occurs after the tower is fully deployed and guyed. Then, all downward forces due to payload, tower section dead weight, guy tension pre-load and guy tension wind loads are transferred by the support frame assembly to the trailer.

Based on this analysis the loads at all significant interfaces were calculated. The interfaces were then proof-tested to load values using a 2X safety factor. During tower tilt-up tests the welds on the strut attachment tabs on the support frame failed. Inspection revealed that the welds had insufficient penetration. New tabs were fabricated, properly welded to the support frame and tested. The L-shaped pivot brackets, fastened with bolts by the manufacturer, were also welded to the support frame by ERIM personnel. The deployment procedure was then revised to position the outriggers in a manner that would allow them to share the tilt-up loads with the strut attachment brackets. This revision changed the outrigger location so that the leg of each outrigger was located forward (toward the trailer hitch) of the support frame and at an approximately 30 degree angle to the axis of the trailer.

Guy Wire Assembly Verification

The guy wire assembly was analyzed for a 75 mph wind speed condition. All welded guy wire attachment lugs located at the top of each tower section were tested with a specially fabricated device that simulated guy wire loads at the welded joint. All guy wire earth anchors were proof-tested with a specially fabricated hydraulic device that tested each installed anchor to a 2.5X safety factor. Calculations were used to verify that the rated cable load of the guy wires was not exceeded. No significant modifications of the guy wire assemblies, other than a change from the duckbill-style anchor recommended by the tower manufacturer and a doubling-up of the replacement auger-style anchors, were required.

5.5.2 Image Processing Results.

The image processing discussion will be based on a 12 second sequence from data collected on June 6, 1995 using LASAR DatacameraTM Serial Number ERIM001. This data is used because:

(1) The 3D sensor started to degrade a few days after this data was collected,

(2) The 3D sensor was pointing to a section of the road which was much closer to the sensor than the required design height of 100 feet. This reduced the maximum operating range, but also reduced the linear road coverage to approximately 120 ft., i.e., 80 feet below the design value of 200 feet. The advantage of this modified operating configuration was that the quality of the raw range data was significantly better than that acquired at the tower height of 100 feet. Section 6.3 will show some data at the VME design range and discuss its properties. The

reduced range data is well suited for the image processing because it meets the data quality assumptions better than the data collected at the increased range.

Figure 5-8 shows raw and processed images from four consecutive even scans (recall that the sensor collects data from both scan directions producing a left-right reversal of apparent image contents.) The first column of images where collected 1.8 seconds into the sequence. The second, third and fourth columns were collected at 2.0, 2.2 and 2.4 seconds into the sequence, respectively. The VME-MS software also used the right-to-left scans collected at times 1.9, 2.1, and 2.3 also, but are not included in this discussion due to the left-right reversals of the images. (The software merges the data after the vehicles have been detected and converted to vehicle features.) The topmost row of images is from the intensity channel. The two lanes of traffic seen on the top side of the imagery are the west-bound traffic on Plymouth road in front of ERIM's building. A large truck and car can be seen moving through the scene in the lower half of the image sequence. An east-bound car can be seen leaving the field-of-view of the t=1.8 frame. A second east-bound car enters the field-of-view on t=2.0 frame. This car can be observed progressing through the scene in the next two frames.

The middle row of images shows the raw range images corresponding to the intensity images shown in the top row. Note that there are two ambiguity boundaries. The first falls approximately between the west-bound and east-bound lanes. The first ambiguity boundary is 143.20 feet from the sensor. The second ambiguity boundary is on the far side of the road at a range of 192.42 feet from the sensor. The truck and the two east-bound cars can easily be seen in this channel. The truck encounters the ambiguity boundary on frame t=2.0 and moves through it in the next two frames.

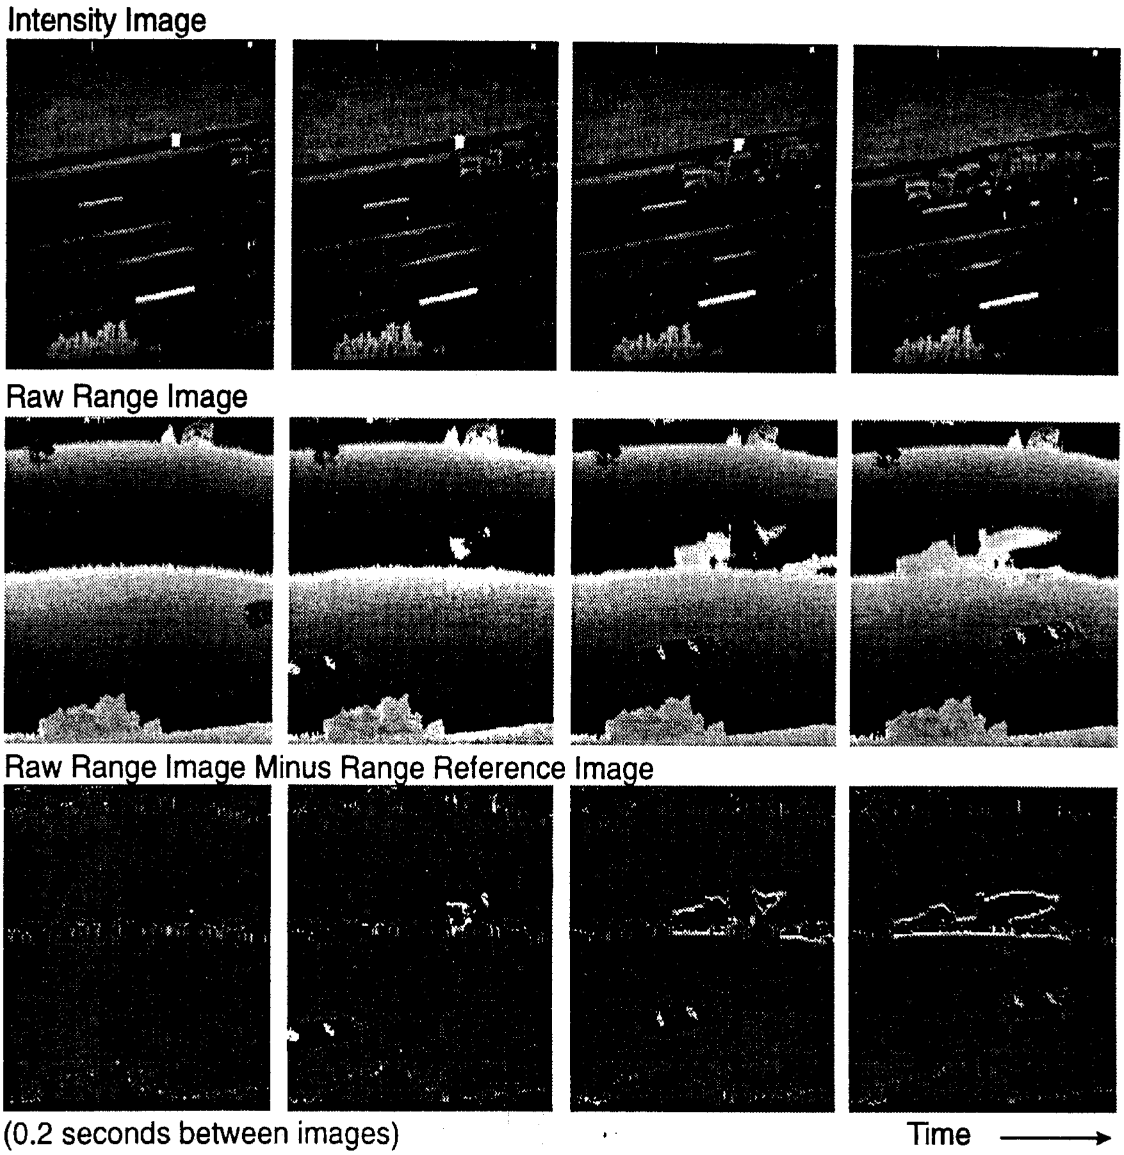

Figure 5-8. Sample Imagery from LASAR DatacameraTM and Processed Results

The bottom row of images in Figure 5-8 are the processed images. These images are the modular differences between the relative range reference image and the raw range images. Figure 5-9 shows the relative range reference image from this sequence of images that was used to produce the processed images. The truck and cars are clearly detectable in the difference images except near the ambiguity boundary. The ambiguity effect shows up in this sequence. If the sensor gave the correct modular range, then the ambiguity boundary would only show up in the range images and not the modular difference images.

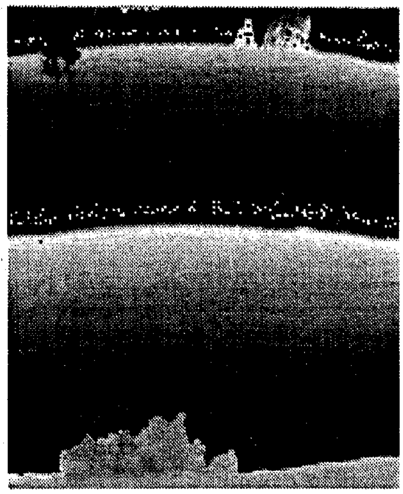

Figure 5-9. Reference Image for Figure 5-8

Figure 5-10 and Figure 5-11, using simulated images, illustrate the expected behavior at the ambiguity boundary and after background subtraction, respectively. Figure 5-10 is a simulated range image with modular range even at the ambiguity boundaries. The ambiguity boundaries are a necessary effect of the 3D technology. Figure 5-11 shows the difference image. Note that there is no ambiguity effect. Because the Perceptron implementation of the 3D scanner gives essentially random numbers near the ambiguity boundaries the boundary effect shows up in the difference shown in Figure 5-8. Except for this ambiguity effect the truck and all vehicles are detectable by the VME detection software. The car near the truck is hard to see in the dithered presentation of the differences, but this car is detected in the first two frames before it becomes lost in the ambiguity boundary.

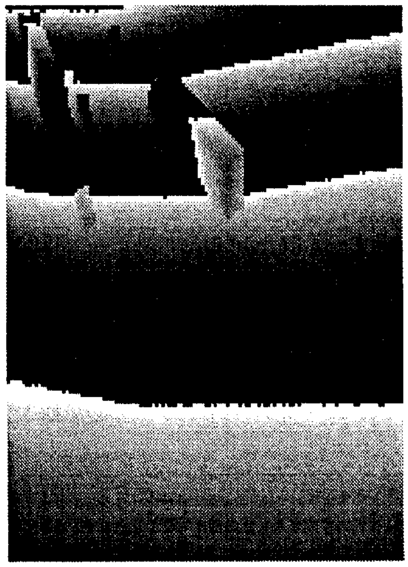

Figure 5-10. Simulated Range Image

These defects in the sensor data meant that vehicle positions and headings could not be accurately determined. Vehicles could be detected and tracked with reasonable accuracy, but that was not sufficient to meet the measurement goals of the VME program. Once the range was increased to operational levels the lack of sufficient laser power resulted in a loss of range data from the road. The sensor started deteriorating (see section 6.3 for example images), and stopped functioning completely during a demonstration. The field demonstration effort was finally stopped at the end of August due to this unreliable sensor performance and the fact that the data from 100 feet was inadequate to support any operational deployment of the VME-MS.

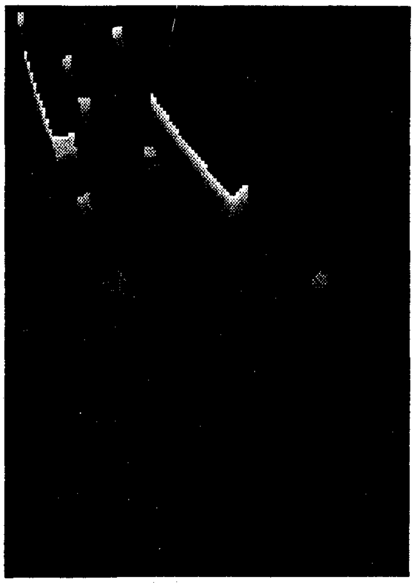

Figure 5-11. Simulated Difference Image

6.0 LASAR DATACAMERA EXPERIENCES

On May 4, 1993, ERIM issued a purchase order to Perceptron for three (3) LASAR DatacamerasTM with an expected delivery on December 23, 1993. Acceptance and delivery of the first two units did not occur until March 30, 1995, more than fifteen (15) months later than expected. ERIM's experiences throughout this procurement are discussed in the following four (4) subsections.

Contractual Summary - The contract is chronologically summarized with particular emphasis on the major causes for the schedule delays. In August of 1994 ERIM relaxed the performance requirements to a minimally acceptable level that was sufficient to support a field demonstration of the concept. The requirements were relaxed so that Perceptron would be encouraged to continue to work on the sensors with a reasonable expectation that ERIM would accept the sensors.

Acceptance Testing - On March 30, 1995, ERIM witnessed the acceptance testing of the three LASAR DatacamerasTM at Perceptron's facilities in Farmington Hills, MI. The results of these tests are briefly summarized.

LASAR DatacameraTM Limitations - The performance limitations of the sensors were the cause for halting the field demonstration test in August 1995. Although lower than required laser power was the fundamental limitation of these sensors, there were several other factors that contributed to their overall poor performance.

Reliability Problems - In addition to the performance shortcomings of the LASAR DatacamerasTM, there were numerous reliability problems with the sensors. Basically they proved to be very unreliable and thus significantly limited ERIM's ability to acquire data to assess the performance of the sensors and the VME-MS as a whole.

6.1 Contractual Summary

On May 4, 1993, ERIM issued a purchase order to Perceptron for three (3) LASAR DatacamerasTM with an expected delivery on December 23, 1993. The specifications for the 3D-laser sensor can be found in Appendix B, VME-MS Design. Acceptance and delivery of the first two units did not occur until March 30, 1995, more than fifteen (15) months later than expected. The third unit was accepted by ERIM on August 24, 1995. The contract is reviewed chronologically from its inception and grouped into four major phases that are characterized by major project milestones.

6.1.1 May 4, 1993 through December 23, 1993

This was the original contract period for the purchase order issued to Perceptron. Although ERIM had been monitoring Perceptron's activities on a regular basis, it was not until October 8, 1993 that Perceptron indicated that they were having technical difficulties related to the procurement of the polygon mirrors. This was less than three months from delivery and Perceptron had not ordered these critical, long-lead components for the Laser Sensor Heads. After ERIM accepted a design change in the polygon that reduced the sampling density in the elevation axis, but simplified the polygon design, the polygon supplier indicated that they could deliver the polygons by mid-January, 1994. Perceptron then asked for a revised delivery date of January 31, 1994. Although this date was unrealistic, ERIM accepted Perceptron's new delivery date.

At this point ERIM contacted the president and CEO of Perceptron, Mr. Dwight Carlson, to discuss our problems directly with him. Mr. Carlson said that he would "look into the matter" and assured us that we had his full support.

Shortly before Christmas, Perceptron informed ERIM that the laser-diode supplier had not delivered any diodes and would not be able to deliver any to the supplier's original specifications. This represented a major performance and schedule impact to the VME Program. At this point in the process, it was not at all clear what Perceptron could deliver, nor when they would be able to deliver three (3) functioning LASAR DatacamerasTM .

6.1.2 December 24, 1993 through August 30, 1994

In a letter to ERIM dated January 6, 1994, Perceptron summarized their development problems which were primarily related to the polygon mirror and laser diodes. Perceptron also established a new delivery date of March 31, 1994 for all three (3) LASAR DatacamerasTM . Neither ERIM nor UMTRI were satisfied with this letter and requested a meeting with Mr. Dwight Carlson. The meeting was held at Perceptron's facilities in Farmington Hills, Michigan on January 31, 1994.

This meeting was the first in a series of meetings that were held throughout the Winter of 1994. In this meeting, as in previous and future meetings, we were assured by Mr. Carlson that Perceptron would make their best effort. Perceptron then proceeded to inform us that the first LASAR DatacameraTM would not be available until April 9, 1994. To help ERIM maintain its schedule for conducting a field demonstration during the late Spring and early Summer of 1994, Perceptron loaned ERIM a 3D-laser sensor that had the same electrical and computer interface as the LASAR DatacameraTM . The availability of this similar 3D-laser sensor did help ERIM with respect to the development and testing of the software for sensor control and data transfer.

After missing the April 9, 1994 delivery of the first LASAR DatacameraTM , Perceptron informed ERIM on May 3, 1994 that the first unit would be ready for shipment on June 11, 1994. Again, representatives from ERIM and UMTRI requested a meeting with Perceptron's management. The meeting was held on May 17, 1994 at Perceptron. Perceptron was asked if they could deliver and what they could do to convince us that they were capable of delivering. Perceptron again assured us that they would deliver and would make a good faith and rigorous effort. They did nothing to convince us that they could deliver. At this point ERIM engaged its corporate lawyer in the negotiation process.

In early June, Perceptron informed ERIM's corporate lawyer that the first unit's delivery date had now slipped to June 28, 1994. On June 29, 1994, after Perceptron had failed to deliver the first unit on June 28, 1994, ERIM made a formal request that

Perceptron re-establish the delivery date. On July 6, 1994, Perceptron informed ERIM that the first unit was ready for acceptance testing and that the remaining two units would be ready for shipment the first week of September. To support an early August Program Review with NHTSA, ERIM temporarily took possession of the first unit and started integrating the LASAR DatacameraTM with the VME-MS. Although the sensor's performance was not satisfactory, ERIM felt that this move was necessary in order to demonstrate the current capabilities of a complete Sensor Station.

After being returned once to Perceptron to fix a severe nodding mirror jitter problem, the first LASAR DatacameraTM was successfully integrated into a Sensor Station and demonstrated to NHTSA at the Program Review on August 5, 1994. At that time, however, it was quite clear that the LASAR DatacameraTM had another major design shortcoming, specifically the lack of any active cooling for the Laser Sensor Head. The lack of active cooling meant that the LASAR DatacameraTM could only be operated in a cooled environment for short periods of time before overheating and shutting itself off. At the Program Review, Perceptron's representative acknowledged that the LASAR DatacamerasTM had to be retro-fitted with active cooling if they were to function in the required outdoor environment. Perceptron agreed to make the necessary changes to incorporate the active cooling and said that two (2) LASAR DatacamerasTM would be ready for acceptance testing during the first week of October, 1994.

On August 30, 1994, ERIM met with Perceptron and relaxed the performance requirements to a minimally acceptable level that was sufficient to support a field demonstration of the VME concept. The requirements were relaxed so that Perceptron would be encouraged to continue to work on the sensors with a reasonable expectation that ERIM would accept the sensors.

6.1.3 August 31, 1994 through March 30, 1995

Throughout September, Perceptron told ERIM that everything was proceeding according to their plan, and on October 10, 1994, ERIM began providing Perceptron with technical support in the testing and evaluation of the first LASAR DatacameraTM. It was immediately apparent, however, that Perceptron was still having problems with Laser Sensor Head cooling and that the sensors were not truly ready for an acceptance test. The lack of a stable thermal environment also affected the optical alignment, thus requiring additional re-work of the Laser Sensor Head. By mid-December Perceptron felt that the first LASAR DatacameraTM was now ready for its acceptance test. ERIM again supported the tests, and again the LASAR DatacameraTM was not meeting the performance specified in the revised acceptance criteria.

On December 22, 1994, Perceptron requested permission to ship the LASAR DatacamerasTM to ERIM on, or before, December 31, 1994. ERIM informed Perceptron that ERIM would not accept the LASAR DatacamerasTM if they were delivered. Perceptron shipped the LASAR DatacamerasTM without ERIM's permission. It took until early February, 1995, before the units were returned to Perceptron, and Perceptron resumed their re-work and testing. By mid-March, acceptance testing was resumed using the VME-MS Sensor Control and Processing Unit to control and capture the sensor data. The results of these tests are summarized in Section 6.2, Acceptance Testing.

6.1.4 March 31, 1994 through August 28, 1995

ERIM's experiences with the LASAR DatacamerasTM after acceptance are described in other sections of this report related to VME-MS development and testing efforts. The terminating event in this very difficult procurement with Perceptron occurred on August 28, 1995 when the LASAR DatacameraTM failed during a demonstration to representatives from NHTSA. The LASAR DatacameraTM was installed in a Sensor Station as part of the Field Demonstration Tests being conducted on ERIM's property facing Plymouth Road in Ann Arbor, Michigan. Later inspections of the Laser Sensor Head (LSH) indicated that something internal to LSH became dislodged and blocked the optical path. This was a new failure mode, but simply another failure that characterized the unreliable nature of the LASAR DatacamerasTM .

6.2 Acceptance Testing

On March 30, 1995 representatives from ERIM witnessed the acceptance testing of the three (3) LASAR DatacamerasTM being developed by Perceptron for the VME Program. As described in the previous section, the Acceptance Criteria were relaxed in attempt to continue the development along a path that would lead to the production of sensors that would have sufficient performance to validate the measurement concept. The revised Acceptance Criteria and the results from the acceptance tests are summarized in the following sections.

6.2.1 Acceptance Criteria (revised)

1. The minimum range performance must be 250 feet. That is the range noise at 250 feet must be no greater than that specified for 350 feet as specified in the following table. The rms range noise values will be obtained from viewing a uniform reflectance target measured at nine (9) points uniformly distributed over the FOV. (At least 90 percent of FOV must exhibit a range noise at, or below, these values.)

RMS Range Noise @ 350 feet	Reflectance
15 in.	5%
10 in.	10%
7 in.	20%

2. The region around any ambiguity point within which the range accuracy could exceed the specified limits established by the range noise must be less than ± 0.5 feet at the ambiguity point. This will be verified by measuring a 20 percent uniform reflectance target centered about each ambiguity point.

3. The reflectance-to-range crosstalk should not increase the RMS range noise by more than 25 percent. This will be verified by measuring a target at constant range with a 20 percent step change in contrast.

4. The response to a step change in range shall be less than 2 pixels. That is, the range measurement at 2 pixels from the step change in range shall have an accuracy limited only by the sensor's range noise. This will be measured at a range of 150 feet when viewing a simple geometric object that has approximately a 2 foot range change from a constant range background.

5. The above values shall not change by more than 1 percent over any 10 minute interval when measured for 5 seconds every minute.

6. The sensor must be capable of measuring a FOV of 30 degrees in the polygon dimension and 50 degrees in the nodding mirror dimension. This will be measured by placing an object with both reflectance and range contrast with the background at the corners of the FOV and verifying that the expected changes appear in the stored data.

7. The intra-frame and frame-to-frame geometric error must be less than 0.5 pixel when measured over a 10 minute interval. Samples will be taken for 5 seconds every minute. This will be measured at a range of 150 feet when viewing a simple rectangular geometric object that has approximately a 2 foot range change from a constant range background.

6.2.2 Acceptance Test Results

Two LASAR DatacamerasTM were accepted (S/Ns ERIM001 and ERIM003) and the third was rejected (S/N ERIM002) due to low-laser power (approximately 100 milliwatts, average, rather than the revised average power of 350 milliwatts, or greater). The tests were conducted with the LASAR Datacameras TM operating at 10 Hz and the VME-MS Sensor Control and Processing Unit was used to control and capture the sensor data.

Except for Criterion No. 2, performance around the ambiguity boundary, the performance of the sensors was evaluated against the remaining six (6) criteria. The evaluation was based on: 1) data that Perceptron had acquired and reduced, and 2) repetition of certain tests to verify that the sensors would repeatably produce similar results. Two of the units satisfied these six (6) criteria. The ambiguity interval boundary test was not conducted because it was Perceptron's position that the phase measurement electronics were working to their design capabilities, and that Perceptron was not in a position to improve the electronics. ERIM was aware that the errors at the ambiguity boundary did not satisfy the revised performance criteria, and was simply left with no choice but to accept the sensors, if the field demonstration was to be conducted.

Although poor performance around the ambiguity boundary did cause problems for the image processing software, this performance limitation was not the cause for the ultimate failure of the field demonstration. The sensor's low signal-to-noise ratio, due to low radiated power, is the primary limitation of the current sensor's performance. When operated at short ranges, such as in ERIM's highbay facility, the sensor's data quality was sufficient to support the intended purpose to accurately measure the position and location of vehicles in the sensor's field-of-view.

6.3 LASAR DatacameraTM Limitations

The 3D sensor did not meet ERIM's specifications with respect to power output, cross talk between the range and intensity channels, and reliability. In addition, the sensor design does not produce modulo range near ambiguity boundaries. Finally, the sensor does not appear to have been calibrated for the distances our project is using.

6.3.1 Power Limitations.

The VME-MS specifications required 1 watt average power. The average laser output is approximately 330 milliwatts. Laser power directly affects the sensor's signal-to-noise ratio. The lower power reduces the effective range well below the expected 320 feet. ERIM believes that there is a problem with Perceptron's range calibration tables when low signal returns are encountered. This belief is supported by the fact that there is a sudden loss of range performance as the return signal approaches an intensity threshold.

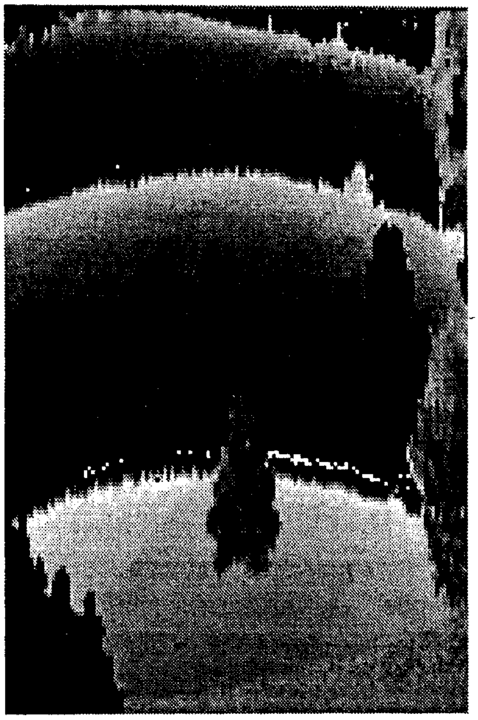

Figure 6-1. Range Image from Top of Tower

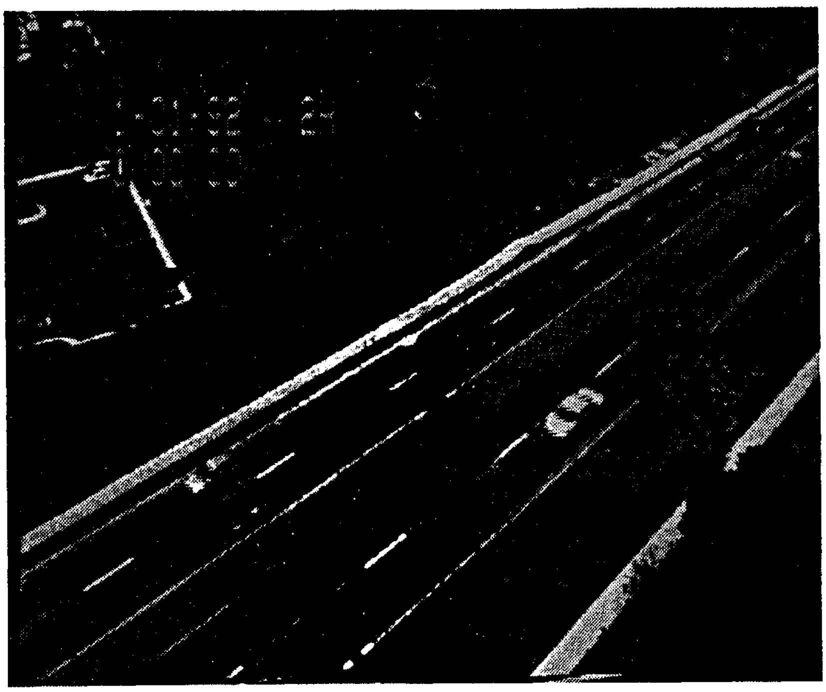

Figure 6-2. Video Image from Top of Tower

Figure 6-1 shows a range image from the 3D sensor. Figure 6-2 shows a conventional visual image of approximately the same section of roadway. The light to dark boundaries, called ambiguity boundaries, are 15 meters apart in range. There is one car in the field of view between the 143.2-foot and 192.42-foot boundaries. These are the first two boundaries, respectively, when counting from the bottom of the image. Thus, the range values before the 192.42-foot boundary are reasonable (except for the ambiguity problem mentioned previously.) The range values above the 192.42-foot transition show a 5-foot to 10'-foot range discontinuity (a massive hole) which is not present in the scene. The range values from the road are not valid above the 192.42-foot line. The low laser power is a major factor in causing this problem, but there are two other contributing factors. The first factor is Perceptron's range calibration look-up table. The look-up table values are not correct for the distances to the road from a 100 foot tower. When the intensity drops below a fixed level the range "correction" from the look-up table introduces a massive error. The second factor is spectral reflection. The aggregate component of the road surface might be reflecting too much of the laser's energy away from the sensor. ERIM has not had this problem before with asphalt viewed from the same angle. Thus, it is believed that the combination of low laser power and incorrect look-up table values are the major causes for these problems. The net result is that the road gives valid range values only when the distance is less than 200 feet and away from the ambiguity boundaries.

6.3.2 Cross-Talk between Range and Intensity Channels.

Figure 6-3 shows an intensity image, associated with the range image shown in Figure 6-1. Note how the intensity channel is contaminated by bright returns along some equal range lines. This effect seems to be independent of ambiguity boundaries. The significance of this effect is that it reduces the ability of the system to use the intensity channel to find reference points within the scene. The reference points are expected to be only a few pixels in size. Note that the noise (anomalous bright returns) occurs in a wide variety of sizes.

Figure 6-3. Crosstalk in Reflectance Image

6.3.3 Ambiguity Boundary Problem.

ERIM has a long history of working with the 3D technology. The 3D sensor has been characterized as a modulo range measurement system. It measures range in the same manner that a watch with no hour hand measures time. It tells you the minutes and seconds, but not the hour. If one has a second time estimate, from other sources, which tells the time to the nearest half hour then the two sources can be merged to give the correct time. This is how ERIM has used the 3D technology. Perceptron's implementation of the 3D sensor returns "random" ranges near the ambiguity boundary. Perceptron now cautions their customers to use only data in the central 70 percent of the interval between the ambiguity boundaries.

The VME-MS requires the acquisition of data across several ambiguity boundaries. ERIM's "fall back" procedure was to detect the ambiguity boundaries by noting the apparent discontinuity in the road (see figure 5-4 for an example of this discontinuity). The position of the ambiguity boundary can be saved in an image mask, so that pixels with invalid range values can be flagged. Vehicles can be detected and tracked between the boundaries. Initially, only vehicles strictly between the boundaries would be tracked. This "fall back" approach loses on the order of 50 percent of the vehicle data, but would allow proof of principal data collections. Later in the processing, more of the data could be used by detecting portions of vehicles crossing the boundaries. For example, if the front of a vehicle is currently in the region with invalid range values, the location of the rear of the vehicle could be detected and tracked. The detection of partial vehicles is a complex task because the regions in the scene where the sensor's values are invalid is a three dimensional phenomena. The top of vehicles enter the transition region before the lower portions of the vehicle. This approach was not pursued due to the basic unreliable performance of the sensors.

6.3.4 Reliability of 3D Sensor.

Only one data set collected in the field met the requirements for algorithm training. Most collected data were contaminated with sensor artifacts such as cross talk between the channels, very low laser power, or excessive noise. The reliability problems are summarized in Section 6.4.

6.4 Reliability Problems

During operation of the VME-MS system, ERIM experienced numerous reliability problems in the Perceptron LASAR Datacamera sensor. This section will outline the history of the three sensors, highlighting the problems with each sensor. Sensor is being used in this section to mean the entire collection of LASAR DatacameraTM equipment, the Laser Sensor Head, the LASAR Sensor Controller, and the computer interface cards.

Sensor #1 (S/N ERIM001)

Sensor #1 was the first sensor to be used by ERIM. After the March 30, 1995, delivery, ERIM used it for laboratory/highbay tests to develop all the sensor control code that resides in the SCPU. This sensor was then used in the initial site deployment tests until it failed due to a electric short inside the Laser Sensor Head. All the sensor data used in this report was produced by this sensor. The sensor was then returned to Perceptron, where the short was repaired, and the sensor returned to ERIM. The sensor then experienced a slow degradation within about a month of use to the point where no usable data was being produced at all. The sensor was then again returned to Perceptron, and it was found that the failure was due to the sensor's receiver in the sensor head. Sensor #1 was returned to Perceptron for repairs.

Sensor #3 (S/N ERIM003)

Sensor #3 was also delivered March 30, 1995, where it remained in storage while sensor #1 was being used for software development and testing. This sensor was then installed in the Sensor Station at the deployment site when sensor #1 experienced its first failure. Sensor #3 failed before any meaningful data could be collected. The failure was associated with a polygon mirror motor controller electronics in the laser sensor head. Upon being returned after being repaired by Perceptron, ERIM tested the sensor, and the data produced from it was deemed to be far inferior to the data produced by sensor #1, hence unusable for the any VME-MS function. No further tests were conducted using sensor #3.

Sensor #2 (S/N ERIM002)

Sensor #2 was the last sensor ERIM accepted, and was delivered on August 24, 1995. This sensor was calibrated specifically for the low-signal returns that were being experienced with the other two sensors. Sensor #2 was immediately installed in the Sensor Station at the deployment site, and failed unexpectedly four days later during a demonstration to NHTSA representatives. Sensor #2 never produced data of a quality equal to that produced by Sensor #1. The sensor was returned to Perceptron for repairs.

7.0 CONSIDERATION OF ALTERNATIVES TO LASER RANGE-IMAGING

The material in this section addresses realistic alternatives to laser range-imaging, the prospects for a video-imaging system in particular, and the technical challenges facing a video system implementation.

7.1 Realistic Alternatives To Laser Range Imaging Sensor

The selection of a sensing technology for the VME-MS is strongly driven by the performance requirement to accurately locate a vehicle along a roadway as a function of time. Table 7-1 shows the basic data requirements of the VME-MS. Motions of vehicles through a roadway segment can be quantitatively characterized by the time histories of these parameters. The data file containing these variables must provide a 10 Hz sampling of the actual vehicle trajectories for monitoring periods on the order of days. The main issues are tradeoff between an active or passive sensor, and between sensor and processing complexity and their related costs.

Table 7-1 Basic Data Requirements of the VME-MS

- The X and Y coordinates of the geometric centroid of each vehicle, relative to a ground-fixed datum (accuracy required ± 6 inch);
- The yaw angle of each vehicle (accuracy required ± 2 degrees at 120 ft. slant range, ±6 degrees at 300 ft. slant range);
- The length and width of a rectangle which outlines the plan-view of each vehicle (accuracy required ± 3 inches)

7.1.1 Sensing Alternatives.

Given that the VME-MS will be a remotely located sensor, and must produce data that is either directly related to range or provide sufficient information for deriving range, some form of range-imaging sensor is implied. Table 7-2 lists the most common active and passive techniques for implementing the basic sensor along with their respective attributes [Best, P., "Range Imaging Sensors", General Motors Research Labs Report No. GMR-6090 (1988)]. Active sensors provide greater control over data quality at the expense of increased sensor cost and complexity. Passive sensor can be relatively inexpensive and simple, but the data processing needed to derive accurate range information can be computationally intense.

Table 7-2. Comparison of Active and Passive Imaging Sensors

CLASS	SENSOR	ADVANTAGES	DISADVANTAGES
Active	3D Laser Imaging	Moderate Cost High Resolution (Range,X,Y,Z) Stable Signatures Req. Moderate to Large FOV Moderate to Low Processing	Eye Safety
Active	Radar (MMW)	Adverse Weather Penetration	High Cost Large Antenna Low Range and X,Y Resolution Processing Intensive Moderate Signal Stability
Active	Range Dispersion	Compact and Inexpensive High Resolution (Relative) Minimal Processing Req.	Eye Safety Experimental
Active	Acoustic Imaging	Inexpensive	Large Antenna Low Resolution Environmental Limitations High Background Clutter
Active	Doppler Imaging	Stable Signature Direct Velocity Measure Moderate Discrimination Min. Background Clutter	Moderate Cost Eye Safety Stationary Vehicles Missed No Range Information
Passive	Visual (i.e., conventional video camera)	Low Cost High Resolution Small/Compact	Ambient Lighting Uncontrolled Contrast Processing Intensive
Passive	Infrared	Moderate Adverse Weather Penetration Variable FOV High Resolution Moderate Cost	Processing Intensive Highly Unstable Signatures Detector Cooling

In principle, an active laser-based 3D imaging sensor appears to be the best candidate for the VME-MS. In practice, our attempt to use such a sensor led to failure due, primarily, to the immaturity of the commercially available 3D imaging sensors. As the first feasible alternative, a passive commercial video camera is recommended for the VME-MS because: 1) reliable passive CCD cameras are commercially available, 2) the higher angular resolution afforded by a visual region video camera is able to provide the required information, 3) the computational burden imposed by image processing and vehicle detection algorithms is manageable, and 4) the video cameras can be readily substituted for the 3D imaging sensors, which are currently installed in the existing VME-MS Sensor Stations.

Among active sensors, as indicated in Table 7-2 above, MMW radar and acoustic sensors require extremely large antennas to meet the positional requirements of the VME-MS. For example, an MMW radar operating at 95 GHz would require an antenna 48 inches in diameter to obtain the needed level of spatial resolution. Acoustic imagers would require even larger antennas.

Among passive sensors, infrared images do provide the capability for imaging at night without artificial illumination. This is the only advantage they possess for this application. Commercial video cameras are at least a factor of ten cheaper than infrared imagers and have much higher spatial resolution.

7.1.2 Data Processing

Perhaps the most challenging aspect of using a passive imaging sensor in the VME-MS is the computational complexity imposed by the image processing and vehicle detection algorithms. As higher resolution infrared imagers become available at reasonable cost, vehicle tracking techniques developed for daylight sensors can be extended to nighttime conditions.

The first challenge in processing video imaging for tracking purposes is to perform vehicle detection. This involves separating vehicles from the background. In a passive system, the appearance of a vehicle depends upon its color. Some vehicles in a scene may appear brighter than the roadway, while other vehicles in the same scene appear darker than the roadway. In the visible spectrum, radiation from the background scene varies with diurnal and seasonal cycles. It also depends on varying cloud, weather and traffic conditions. Operation in the visible band also introduces shadowing which in turn produces additional clutter that further complicates the data processing algorithms. Vehicle detection by change detection is recommended for the VME-MS because: 1) both bright and dark vehicles can be detected, 2) by using a CCD camera with a variable exposure time some of the illumination variability can be ameliorated, and 3) by continuously adapting the reference scene (by median filtering for example) the localized effects of clouds and cloud shadows can be largely nullified. The contrast being exploited by change detection in passive visible ban imagery is the contrast in appearance between and individual vehicle and the surrounding roadway.

The second challenge is processing video imagery for location detection purposes is to perform range determination. Using stereo vision is beyond the scope of the proposed effort. Figure 7-1 illustrates the classic problem of a single camera system - the apparent location of a vehicle is influenced by its height. However, this problem is not insurmountable. Range determination based upon the lowest point on a vehicle is recommended for the VME-MS because: 1) using a surveyed roadway allows precise range determination for the scene when no vehicles are present, 2) the observed position where a vehicles tires contact the roadway provides location on that datum, and 3) knowledge of that range allows the extent of the vehicle in the (angle, angle) coordinates of the camera to be transformed into vehicle dimensions of length, width, and height.

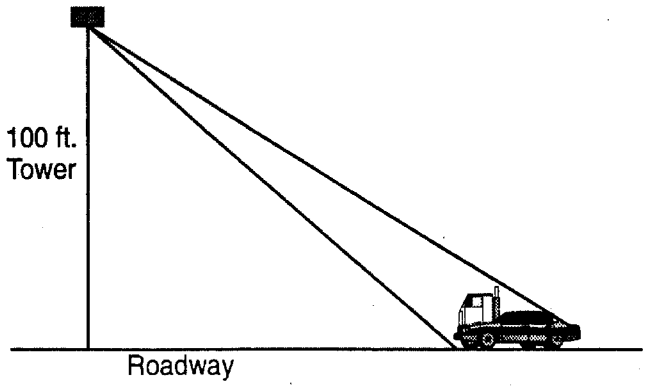

Figure 7-1. Height effect on location

The third challenge is in processing video imaging for tracking purposes is to keep up with the data rate of the video images. The finer angular resolution of the CCD video camera, 484 by 768 pixels, requires more computational throughput to stay with the 10 Hz sampling of the roadway environment. Switching from Intel 486 processors to Intel Pentium^{TM} processors is recommended for the VME-MS because: 1) this provides significantly increased processing speeds, and 2) the PCI bus on the Pentium is needed for the digital video data from the CCD camera.

Aside from the phenomenology on which it is based (i.e., visual appearance versus relative range), the CCD video camera system will be transparent to the remainder of the vehicle measurement equipment. Vehicle locations will be determined 10 times per second as they traverse the 60-foot x 200-foot portion of a roadway viewed by a single sensor. Observations on vehicles from two adjacent sensors will be merged, over time, to produce track files for individual vehicles as they transit a longer segment of the roadway.

Table 7-3 lists the performance strengths and weaknesses of the CCD video camera-based system with respect to the original laser range imaging system. If the available laser imaging sensors had worked, we would have used them. Given the immature level of the available laser sensor, the CCD video camera represents a viable alternative. Albeit an alternative that exchanges sensor risk for the challenge of software and algorithm development.

Table 7-3. Performance Comparison

	CCD Video Camera	Laser Range Imaging Sensor
Phenomenology	Measures reflected light from each scene element (external lighting required)	Measures relative range to each scene element (no external lighting needed)
Roadway Location of Objects	Inferred from observed (height and shadow complexity) position in a surveyed scene	X,Y, Z Determined directly from sensor data
Scene Observation Rate	10 frames/second	10 frames/second
Image Data Volume & Rate	484x768 8-bit pixels, 3.72 Megabytes/sec	128x180 2x16-bit pixels, 0.92 Megabytes/sec

7.2 VME Video Based System.

The following sections will describe the modifications to the VME-Measurement System (VME-MS) to change it from a laser-based to a video-based system. The installation of new hardware and software elements will be very straightforward due to the modular design of the system. The system architecture will not be changed and the video-based VME-MS will maximize the use of existing hardware and software elements. The significant changes are:

Hardware:

- Replacing the LASAR Datacamera™ with a Pulnix digital CCD video camera
- Replacing Perceptron's computer interface with a Video capture board (frame grabber)
- Upgrading the Sensor Control and Processing Unit to handle increased processing and video capture requirements

Note: Any new hardware necessary to execute this demonstration will be loaned to the Program by ERIM.

Software:

- vehicle detection and feature extraction modules that reside in the Sensor Control and Processing Unit
- control modules to manage real-time / non-real-time execution

Note: The message-based control software and track-file management software are not changed.

The modifications to the current laser radar-based VME Measurement System will be highlighted in each section.

7.2.1 Current Hardware System

The current system configuration is shown in Figure 7-2. The system consists of multiple Sensors Stations, each consisting of a standard video surveillance camera and 3D laser sensor mounted on a tower, and an Electronics Unit at the base of the tower. Each station's Electronics Unit contains a commercial off-the-shelf Video Recorder to archive the video data for later viewing and a Sensor Control and Processing Unit (SCPU). The SCPU controls the 3-D laser sensor's data acquisition and processes the 3-D image data. The processing results are then passed to the Master Sensor Station.

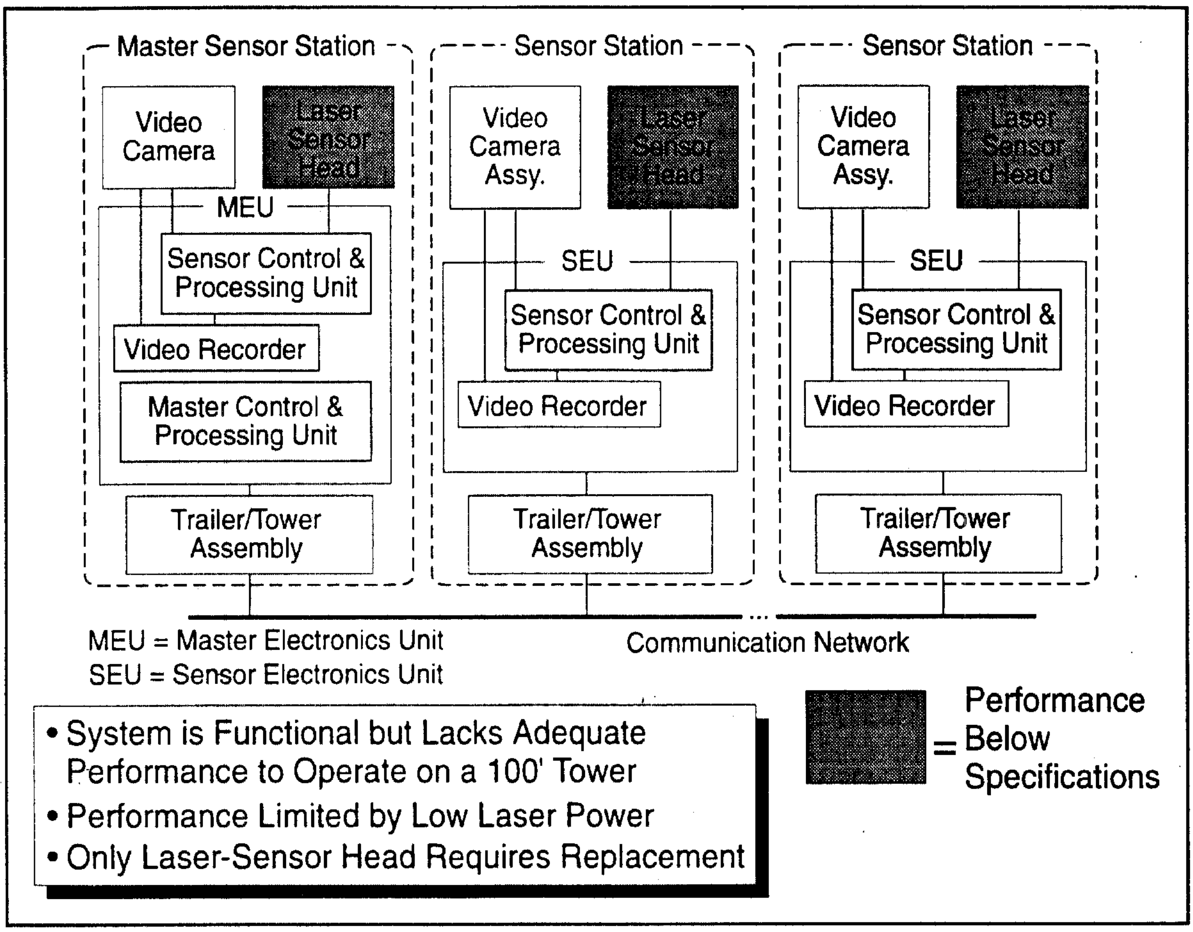

Figure 7-2. Laser-Based VME-MS Block Diagram

The Master Sensor Station is like the other Sensor Stations, except it hosts the Master Control and Processing Unit (MCPU). This unit merges the information of all the sensor systems into a unified database which is archived using a 4 gigabyte data-compressing DAT drive. The MCPU does not require any modifications as a result of changing the type of sensor.

The shaded regions of Figure 7-2 highlight the Laser Sensor Heads, which are the sole item of the VME-MS that are performing below specifications, making the system incapable of producing useful track file data.

The following three sections highlight the modifications to the VME-MS that are required for a video-based implementation. Section 7.2.2 provides an operational overview following by descriptions of the hardware and software changes in Sections 7.2.3 and 7.2.4, respectively.

7.2.2 Operational Perspective of Video-Based Approach for VME-MS.

It is informative to walk through the major VME-MS functions from an operational perspective and explicitly describe what portions will be altered and what has been retained. This section will do that in a systematic manner that begins with the viewing geometry of the sensor, moves on to the image and data processing and then covers the retained functionality.

Viewing Geometry and Phenomenology

The same towers which were deployed for the individual Sensor Stations in the original VME-MS will be retained for the video-based approach. Deployment in the fully extended (i.e., 100-foot high) mode will be used to minimize vehicle-to-vehicle masking while monitoring 200-foot section of the roadway just as it was when laser range imaging sensors were being employed by the VME-MS. Here again, traffic information will still be collected at 10 Hz.

The video based approach uses ambient illumination (either natural light or light from street lamps) as opposed to the active laser illumination used by the laser range imaging sensors. Generally speaking, natural light can vary by about two orders of magnitude from a bright clear day to a heavy overcast day with light rain. We expect, however, to use a software controlled feedback technique to vary the exposure time so that the digital imagery from our passive sensor remains relatively constant in brightness.

The video based approach will use a 484 line by 768 pixel intensity image as opposed to the 128 line by 180 pixel range image and intensity image pair of the laser range imaging sensor. The higher spatial sampling rate of the video-based system is needed to accurately recover the local roadway coordinates of vehicles.

Image and Data Processing

As illustrated in Figure 7-3, the same software structure and the same operating system that was employed in the original VME-MS will be retained for the video-based approach. In addition, the existing message passing software will be available to ease the task of determining sensor-to-sensor correspondence between vehicles.

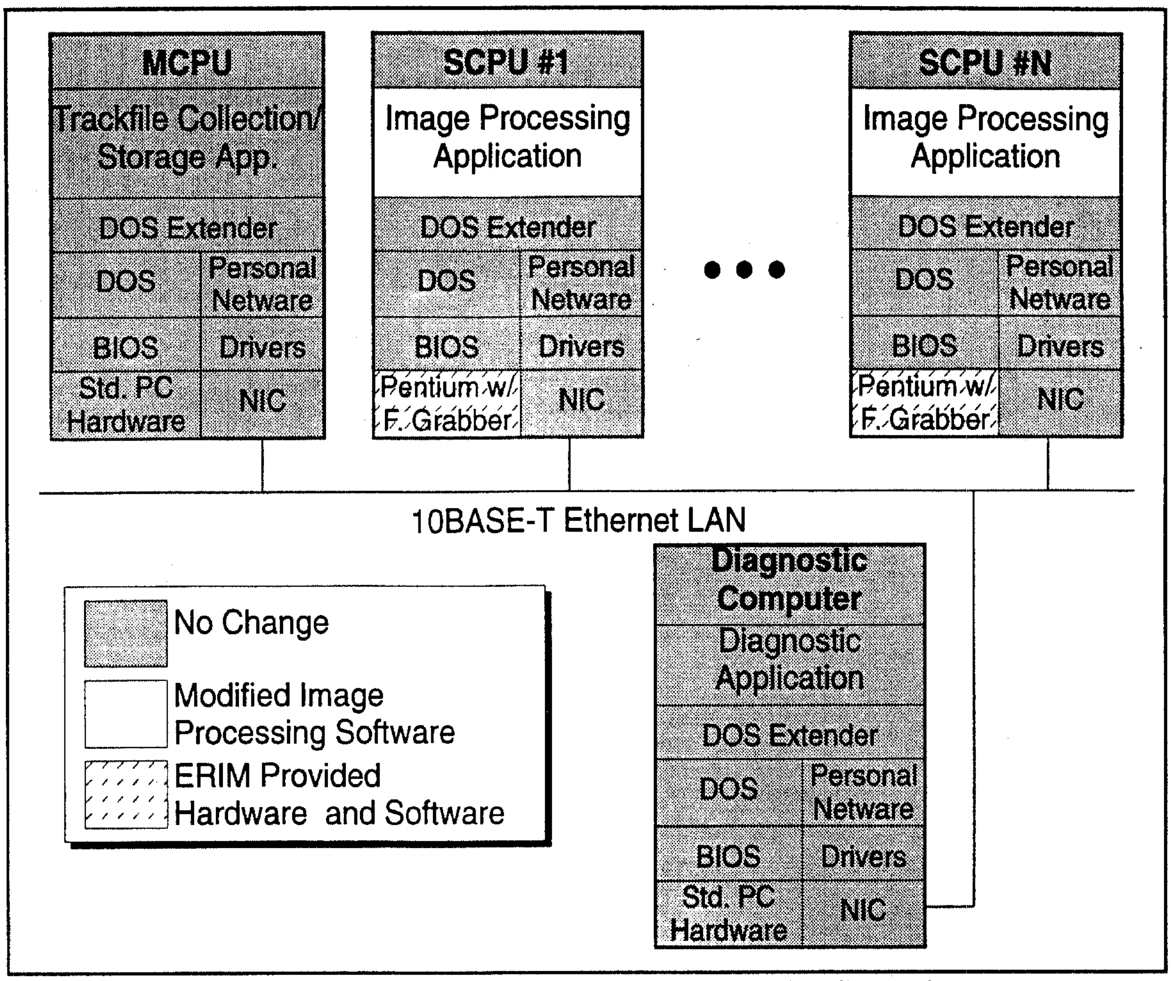

Figure 7-3. Video-based VME-MS System Configuration.

The video-based approach will employ a video frame grabber (i.e., an analog-to-digital converter) installed on a PCI bus whereas a vendor specific digital interface was used by the laser range imaging sensor. The PCI bus on an Intel Pentium computer will replace the Intel 486 VME-bus single-board computer used in the sensor stations.

The largest difference between the video-based approach and the original VME-MS will be in the vehicle detection and feature extraction portions of the image processing algorithms. The video-based approach will employ change detection based upon the visual contrast between a vehicle and the roadway to find the vehicles. In the laser range imaging sensor the change detection employed was based on the physical height of a vehicle relative to the empty roadway. In addition, the video-based approach will infer the local roadway coordinates of vehicles from the apparent position where they touch the surveyed roadway. In the laser range imaging sensor, measured range values to the nearest two sides of each vehicle were used to directly compute the location of the vehicle's centroid relative to a ground fixed datum.

The video-based approach will produce intermediate state-labeled images that may be efficiently stored using the run length encoding capabilities of the DAT drive, thus permitting data collection to proceed at real time rates. Feature extraction and track file production from these state-labeled images could then be done off-line. This feature has been incorporated into the software design because it may not be possible to perform all of the image processing tasks in real time.

Retained Software Control Functionality

Aside from the sensor phenomenology and vehicle "features" employed, the software control functionality of the VME-MS will be retained by the video-based system. Figure 7-4 illustrates the functionality of the individual sensor control and processing units (SCPU) as well as the control functionality of the Master Control and Processing Unit (MCPU). This diagram is identical whether a video-based approach is taken for the VME-MS or whether a laser range imaging sensor is employed. Global track files will be produced in the same manner and format; no changes are necessary.

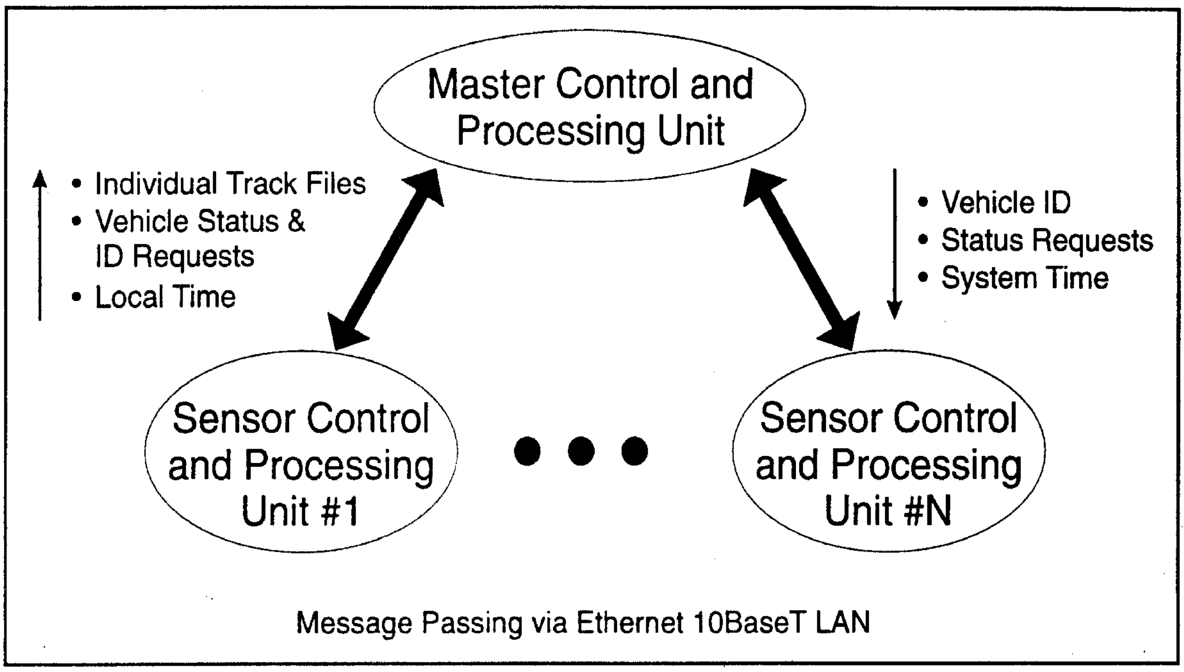

Figure 7-4. System Functionality.

7.2.3 Hardware Modifications

To convert the existing system into a video-based system, the 3D laser sensor will be replaced with a Pulnix TM-9701 digital CCD camera, the associated data acquisition hardware will also be replaced with a digital frame grabber, and the SCPU computer will be replaced with a higher-performance Pentium processor. Figure 7-5 shows the block diagram of this new system configuration. Figure 7-6 and Figure 7-7 are new resulting Sensor Station and Master Sensor Station configurations, respectively.

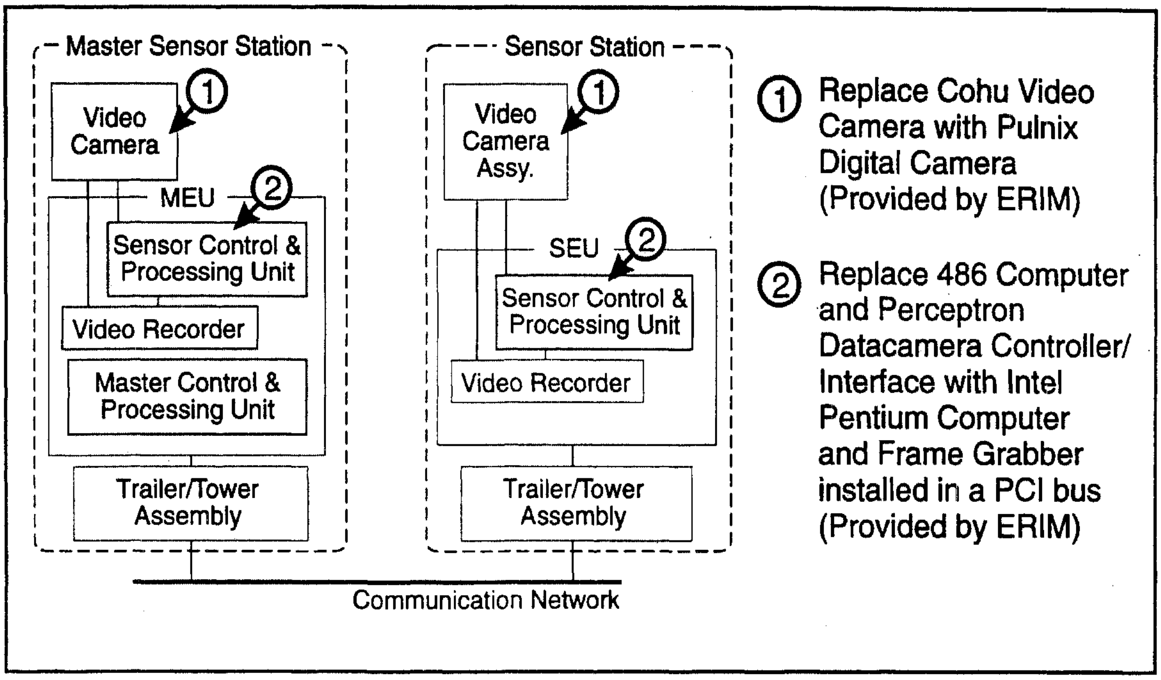

Figure 7-5. Changes to Block Diagram to Convert System to Video-Based.

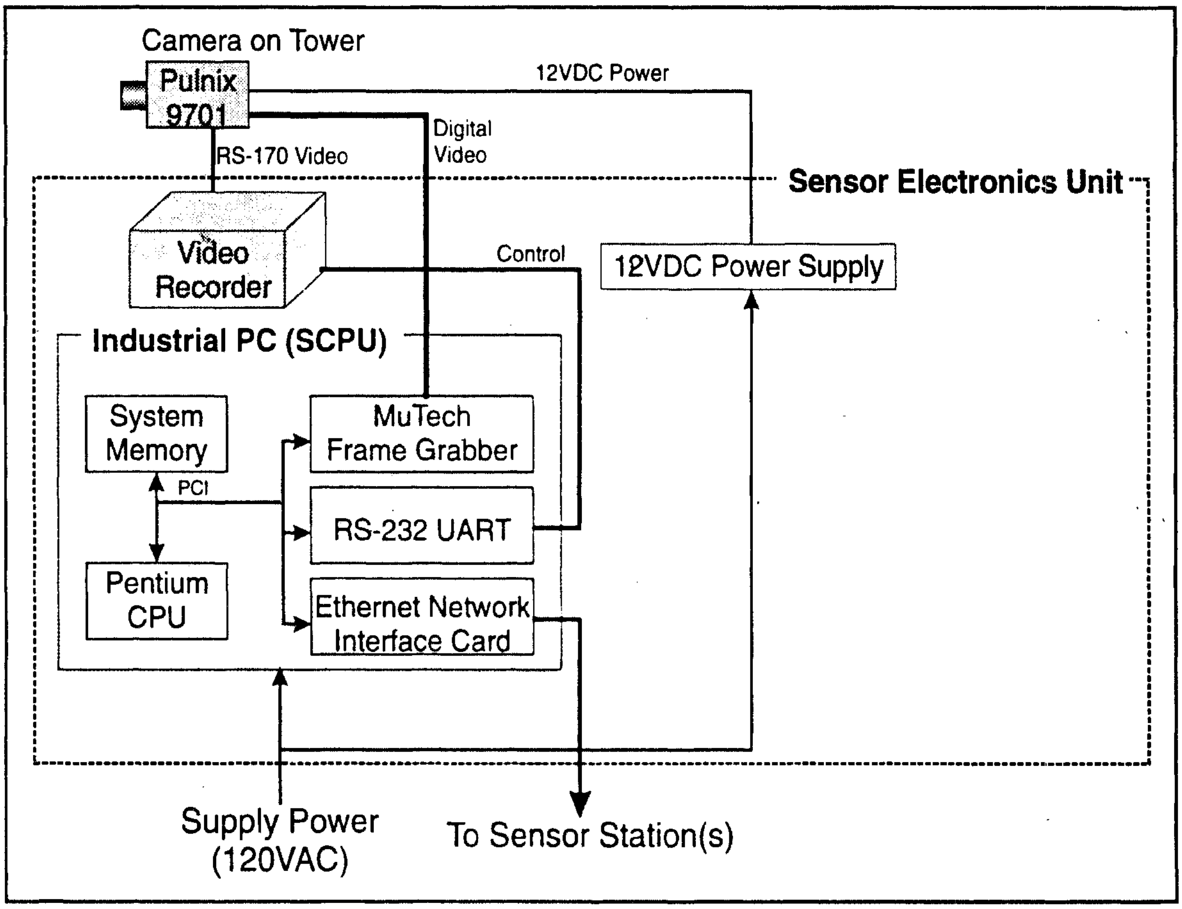

Figure 7-6. Video-based Sensor Station Configuration.

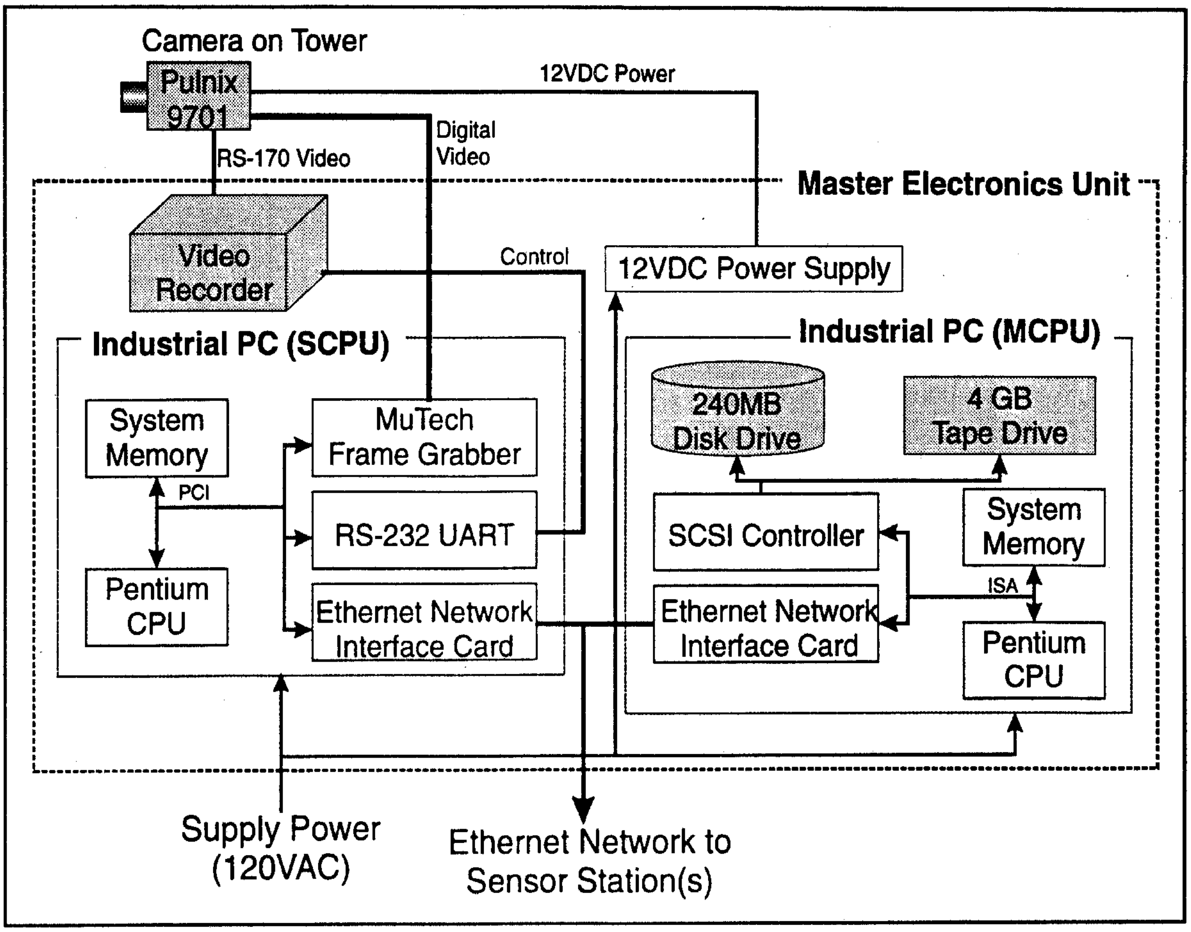

Figure 7-7. Video-based Master Sensor Station Configuration.

Tower & Trailer Assembly

The video-based system will still require the same geometry as the laser radar-based system. The design of the tower and trailer will not need any modifications other than new mounting interface for the digital camera that will replace the laser-radar sensor.

Sensor Electronics Unit

The video-based system will not use the Perceptron laser radar sensor, nor will it use the VME-bus computer and the 3D interface electronics. The laser radar sensor will be replaced with a Pulnix TM-9701 digital CCD camera. The SCPU computer and 3D interface electronics unit will be replaced with a commercial off-the-shelf Intel Pentium rack-mounted computer with a commercial PCI-bus frame grabber installed. The existing software designed for the old SCPU will be used in the new SCPU computer. The Pulnix cameras, PCI-bus frame grabbers, Pentium computer and associated hardware for two Sensor Stations will be loaned to the Program by ERIM for the duration of the demonstration.

Master Electronics Unit

The Master Control and Processing Computer will be unchanged from the previous design. The only changes of the Master Electronics unit are in the components that the subsystem have in common with the Sensor Electronic unit.

Grounded Bulkhead

The grounded bulkhead will be modified to accommodate the new sensor. Most of the surge protectors for the Perceptron Datacamera^{TM} that are located in the bulkhead will be retained for use in the interface with the digital camera. The digital camera will require some additional surge protection since the number of digital lines and type of digital lines varies from the Datacamera.

The surge protection for the power and communication network will remain unchanged.

7.2.4 Software Modifications

The software modifications required for the video-based system are mostly in the image processing algorithms, since the type of data from the CCD camera differs from the laser-based Perceptron Datacamera imagery. Since the sensor hardware will be replaced, the associated control software also changes.

The next two sub-sections will explain the required changes to the control and image processing software.

Control Software

Two parts of the control software will have to change. First is the sensor control, since the sensor interface has changed. Second is the real-time control, since the image processing might not be all real-time.

The sensor control software for the video-based system will use a frame grabber on the SCPU's PCI bus to control the camera instead of sending commands through the RS-232 port to the 3D laser controller. This software already exists for the camera and frame grabber to be used, and since it was developed for the same software environment as the laser-based VME-MS system, it can be easily integrated back into the VME-MS software. This software has provisions for all the required camera control, and has been tested extensively with the frame grabber that will be used. The only modification of the code that is desired is the implementation of continuous (closed-loop) exposure control. Existing software exposure control is open-loop.

Since the image processing might have to be split into a real-time and non-real-time part, the control software that handles the information exchange between software modules will have to be changed. This change is minor, since the software will change the destination of the real-time image and data processing output to a storage device. This information will then be read by the control software and passed to the non-real-time portion of the image and data processing software at a later time.

Image and Data Processing

Figure 7-8 shows the logical image and data processing steps that are common to the laser- and video-based systems. The image from the sensor is processed at each sensor station for vehicle detection, feature extraction, inter-frame correspondence, local coordinate transformation, and time-tagging of position and heading information. This position and heading information is then passed to the Master Sensor Station for processing into one global vehicle track file.

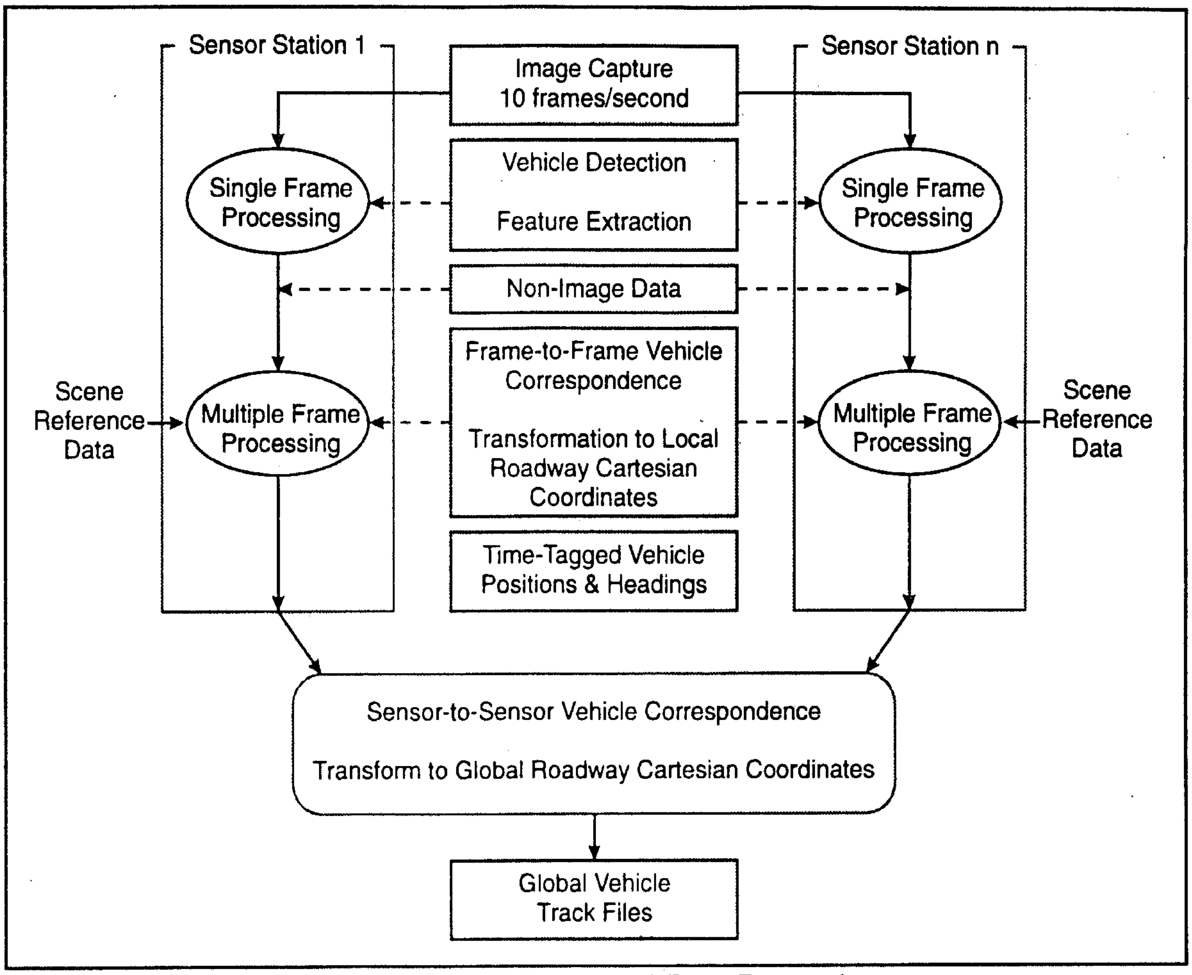

Figure 7-8. The Image and Data Processing.

The software modifications required can be found in the details of how the above logical structure is implemented.

The video-based processing might be split into two main sections: image processing and feature extraction. There is a significant chance the processing time requirements will violate real-time constraints. This is because:

- The replacement sensor has a much greater angular resolution resulting in a significant increase in the amount of data to process.
- The image data no longer supplies range directly, so the position calculations are more complex, thus requiring more computations.
- Finally, since the video sensor is passive, vehicle detection is also more complex due to external illumination and associated variations.

Figure 7-9 shows the detailed processing steps for the image processing and feature extraction. The shaded components will only be required if the image processing cannot be performed in real-time. Section 7.3.3 discusses the two step image processing alternative to a complete real-time processing goal.

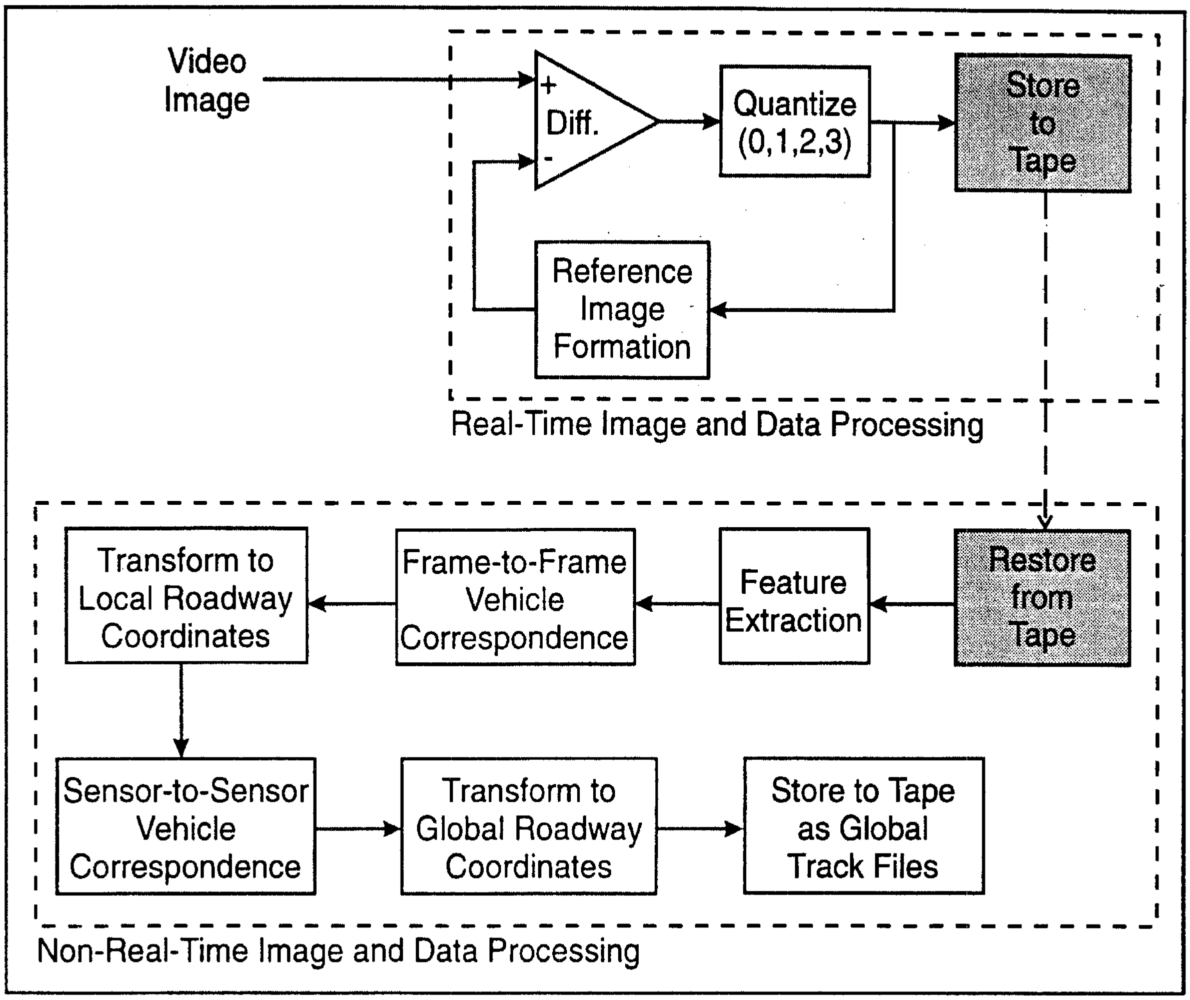

Figure 7-9. The Image and Data Processing of the Video-Based VME-MS.

7.3 Technical Challenges In Applying Video Sensing.

Given that the sensing hardware will now be a passive imaging CCD camera and that this is proven technology, the technical challenges have been shifted from the hardware domain to the software domain. The software elements can be coarsely divided into image processing, data processing and real-time control. The effects on the software design of changing the sensor will primarily impact the image processing element. The nature of these impacts and ERIM's approaches for addressing them are now discussed.

The image processing component must convert millions of bits of image data into vehicle related features suitable for extraction of the position, length and width of the vehicle at the time that the vehicle was first detected. The change from a sensor which directly interrogates the geometric nature of the environment to a sensor that receives reflected visible light has a direct and significant effect on the image processing components. Our approach is to have the image processing element compute the same outputs independent of the sensor. This will localize the effect of changing sensors.

The image processing approach is still based on contrast from an estimate of the road without vehicles, however contrast has a much different meaning with the passive approach. The fact that the contrast can no longer be measured in inches illustrates the changes that must take place within the image processing component. Contrast now means that a region no longer looks like the road did a short time ago. This change might be because a vehicle is blocking the camera's view of the road in which case we have detected a vehicle. Another possibility is that the road now looks different because a cloud has just moved in front of the sun and there are other reasons for contrast changes.

Once a vehicle has been detected the information is still in the form of an image. The 3D laser's information, at the same stage of the image processing, was the vehicle's position in polar coordinates. The new information is in angle-angle form without range. The image processing must derive the x-y-z information indirectly. These two effects, high variability of the vehicle-to-roadway contrast and determination of the vehicle's position, represent the major impacts of changing to a passive sensor. Although the intent is to localize the effects to the image processing portion of the software, there could be impacts on the data processing and real-time execution.

If the new image processing component produces the same features as the 3D laser-based system, then the data processing will remain mostly unchanged. The Sensor Station data processing component will be doing the vehicle tracking so information from the current image frame is added to the correct vehicle file. The Sensor Station computer must send track parameters and later, when the vehicle leaves its FOV, it must send the track file to the Master Control and Processing Unit (MCPU) in the Master Sensor Station. The MCPU must forward the track parameters to the correct station and when the vehicle has left the scene, it must merge all of the station track files into one global track file. The impact on the data processing software from changing to a passive sensor are expected to be minor. The most likely change is that the track files may be produced off-line due to limited real-time processing resources.

There is a significant chance that the image processing time requirements will exceed the time available. There is a natural place in the image processing chain to partition the image processing into two parts. The first part contains those components that must execute in the full image environment. The second part contains those component that can function in a reduced image environment or a symbolic environment.

The first collection of components must execute while the data is being collected. The data size must be reduced by this first phase so that the transformed data can be recorded within the time constraints imposed by our recording equipment. The second set can be executed in a second pass from the recorded data.

7.3.1 Image Processing Software.

The image processing must convert information in the form of pixels into position, length and width vehicle parameters. The Sensor Station data processing component must take this data from one image and update a vehicle track file. The information must be added to the correct vehicle file. The current position and time are added to the file while the vehicle length and width estimates are updated using the new and old estimates. Tracking information that would be useful for the next station must be sent (via the MCPU) before the vehicle is seen by that sensor. When the vehicle has left the station's FOV the vehicle's track file must be sent to the MCPU.

Our approach is to keep the non-image processing components in their current form by requiring the new image processing components to produce the same features as before. There are several potential problems with replacing the laser with a passive sensor. Both the problems and our approach are now discussed.

The video camera uses visible reflected light. This means that the system needs light from the sun or a well lighted section of road. Our initial approach is to collect data only during daylight hours. The question of using the system at night, where there are artificial lights, can be addressed at a later time. Even in the daytime there are changes in the illumination which have nothing to do with the changes in the vehicle/roadway environment. One benign change is the sun's position. Over a short period of time the change is small but over large fractions of an hour changes in shadowing and illumination can be significant. A cloud moving so that it blocks (or so that it no longer blocks) the sun will cause a major change to which the system must adjust.

Our approach is to handle varying illumination at two levels. The first level exploits the computer's capability to vary the camera's exposure. The approach is to detect changes in illumination by having the computer monitor places in the FOV that are not expected to be blocked by vehicles. Statistics from these regions can be used to detect illumination changes. If the mean of these regions is decreasing (or increasing) the exposure time can be increased (or decreased) to compensate for this change. The second level in our approach is to maintain an estimate of the road image when there are no or few vehicles present. This is done by using a median tracker on a pixel-by-pixel basis. When traffic is light and moving, any one pixel will not be blocked by vehicles most of the time and the median tracker will ignore the vehicles as short time duration transients. This median image is called the reference image. When traffic is not moving, or is heavy, the tracker will be restricted to those regions that have no vehicle detections. This approach for estimating the reference image means that the system must be initialized when traffic is light and moving.

Vehicle detection is now based on the assumption that the change in a region is due to a vehicle blocking the road from the camera and the road shaded by the vehicle. Shadows are both a source of information and confusion to passive vision systems. If the entire changed region (both vehicle and shadow) is used as the region where the vehicle is blocking the road from the camera's view, then the position and size estimates will have a large errors.

Our approach is to estimate the shadowed road's statistical signature. This can be done by using knowledge of the sun position to partition the negative contrast region into dark vehicles and vehicle shadows on the road. If there are objects such as telephone poles which cast a shadow across the road, then this information can also be used in obtaining a shaded road estimate. At high sun angles the shadows will be directly underneath the vehicle, giving the system a good estimate of location of the vehicle on the road. At other times the shadow will appear to the side, in front or in back of the vehicle. The knowledge of the sun's position will allow the algorithm to know where to look within the changed region for shadow and estimate its extent. At this stage in the processing, we have a general estimate of the shaded road signature and some pixels within the negative contrast region which are shaded road for that vehicle. This knowledge is used to partition the negative contrast region of the vehicle into "vehicle with negative contrast" and "vehicle shadow." This is the last image processing step that is performed on the entire image. All further image processing is performed on image subsets containing the detected vehicle and its shadow. This stage of image processing is of special interest because it is a natural break point to divide the processing load between real-time and off-line. Section 7.3.3 discusses the division of the process in the time domain.

7.3.2 Data Processing Software.

Once the vehicle parameters (position, length, width) have been estimated, the processing components remain the same. The station data processing component which predicts where the vehicle will appear in the next frame will still track vehicles within one sensor's FOV and between sensors' FOVs. The only change in the data processing software is the added flexibility of having two possible methods for producing track files. The system will produce the track files in real time or in a two step process. The reasons for this flexibility requirement and how the system design has been adjusted to handle it are discussed in the next section.

7.3.3 Real-Time Processing

The goal is to produce track files in real time. The image processing components have not been defined at this time. Thus, meaningful run-time estimates cannot be made. Time estimates will not be available until the form of the image processing algorithms are known at a very detailed level. The general form of the image processing is known well enough to place bounds on the run time. The upper bound on the run time is too high for all of the track file computations to be performed in real time. The lower bound does meet the real-time constraints. If the run-time characteristics of the image processing are fast enough, then the vehicle track files will be produced and saved on tape in real time. If, on the other hand, the image processing executes too slowly, then the image will be reduced from an image with 256 possible states per pixel into a reduced image with four possible states per pixel image. These states would include "no significant change from reference," "significantly brighter than reference," "significantly darker than reference," and "vehicle shadow." This feature image will be produced under both scenarios. If features can be extracted from this reduced image so that the total time is within 0.1

seconds, then the feature extraction will take place while the data is being collected. If features cannot be extracted within the time budget then the reduced image will be compressed using "run length encoding" and written to tape. Figure 7-9 shows the data flow in this two staged image processing.

Most of the software is common for these two environments. In the real-time scenario, there are 3 (n+1) groups of processes sending messages, where n is the number of towers. In the delayed processing mode, the processing is broken up into two parts. The first is the data collection part where there are 3 groups of processes each recording feature images. The second part has 3 (n+1) groups of processes residing on one real computer. From a software point of view the only difference is the image processing is not producing the reduced image but reading it from tape and decompressing it. The other components are the same. The communications among SCPU1, SCPU2, and MCPU do not know that their messages are going to a process on the same CPU. One might say that, in the second part of the delayed processing mode, we are simulating the multi-processor and multi-tower environment of the real-time system.

8.0 CONCLUSIONS AND RECOMMENDATIONS

Insofar as the VME project has only designed and constructed a measurement and processing system, no conclusions or recommendations have been generated that deal with observations of the vehicle motion environment, itself. Nevertheless, conclusions have been made pertaining to the state of the design, construction, and evaluation of the VME measurement and processing system, as cited below:

The current state of the industrial art will not support application of laser range-imaging as the VME sensor.

Mechanical and electrical package is otherwise ready, providing:

- portability (via trailer apparatus)
- extendible platform up to 100 feet elevation
- temperature-controlled enclosure
- lightning protection
- analog CCD surveillance equipment
- processing features that include:
  - conversion of raw track file data to VME database format
  - searching and organization of VME database information / data
  - text, graphic, and animation display of VME database files
  - Kalman filtering
  - crash detection calculations
  - inter-vehicular range and angle-of-attack computations
  - export of various data for access by commercial programs

Some limitations remain regarding mechanical strength vis-a-vis wind loading.

Software for establishing "correspondence," which splices together one continuous track file for each vehicle that crosses the boundaries of adjacent sensor stations, has been proven using simulated data.

Processing system has been shown to properly execute the previously cited features using simulated track files.

Kalman filter calculations effectively estimate the time histories of ancillary variables, from raw track file data, including driver steering and braking / throttle control activity.

Sample data has been obtained showing the high resolution capability of digital CCD imaging. (And, in a parallel project, a system for controlling the digital CCD camera and processing its images at high frame rates has also been demonstrated.)