Automated Volume Estimation of Haul-Truck Loads

Elliot Duff
CSIRO Manufacturing Science and Technology
Queensland Centre for Advanced Technologies
PO BOX 883, Kenmore Qld. 4069 Australia
HTTP://WWW.CAT.CSIRO.AU/CMST/AUTOMATION

Abstract

This paper describes technology (hardware and software) that has been developed to estimate the volume of material contained within the tray of various haul trucks. Such information is critical to truck utilization and thus can provide information that could lead to significant gains in productivity. The system consists of a field computer and two proximity lasers. The lasers are placed high above the roadway to enable them to scan each truck as it passes underneath. It is now technically feasible to develop a system that can operate on-line, where the volume can be transmitted to the truck, or to mine management, as it is being scanned. Such immediate feedback provides important data to the optimization of automated dispatching systems.

1 Introduction

Haul trucks are large and costly surface machines used in open-pit mining to remove large volumes of dirt (overburden/ore). Although the haulage of dirt may seem a simple exercise, there are a number of design and scheduling issues that need to be resolved to optimise the performance of these trucks. In particular, whilst trucks are designed to carry *weight*, they are used to move a *volume* of dirt. This difference in function is natural. Trucks are machines that are designed to withstand mechanical load, thus maintenance engineers need to know how much weight the vehicle will carry. The utilization of each truck can be improved by maximizing the volume of dirt that the tray can carry without overloading the truck. In the past, the volume carried by the truck has been inferred from a volumetric survey of the mine (divided by the number of loads) or a weight survey of each truck (bulk density of the material divided by weight). Neither technique actually measures the load in the tray and does not take into account changes in packing, or swelling, that the load undergoes in the tray. With the development of laser scanning technology it is now possible to measure the in-situ volume of each load directly.

2 Hardware

The hardware for this system consists of a robust field computer and two scanning lasers (see Figure 1). The lasers are inexpensive, eye-safe proximity laser scanners (PLS) from Sick, Germany. In our experience, they have proved themselves to be robust, reliable and accurate sensors. They consist of a low powered laser that is reflected off a spinning mirror. By measuring the time-of-flight of the return signal, it is possible to measure the range (up to 50 m) of objects in its field of view (180^o). They have a resolution of 5cm in $1/2^o$ increments at a rate of 25Hz. Communication with the field computer is via an RS422 cable. The two lasers are mounted orthogonally, approximately 5m above the truck (see Figure 2). This distance is a compromise between angular resolution and the effects of shadow.

The field computer consists of a ruggedized PC (Teknor single board computer) with a Pentium 200MHz processor running LynxOS. A dedicated embedded communications system is used to achieve high speed communication with both lasers (500KBaud). The embedded system (see Figure 3) consists of two components: on the left, an IPcomm360 which is a commercial IP module (Industry Pack) that uses the Motorola 68360 chip; on the right, a PC104 IP carrier which has been designed and built by the Automation Group. The choice of PC104 and IP has been historical. Several years ago, when we were using VME systems, IP modules were chosen because they were designed to be architecture independent. Two years ago, when we needed to move to a more mobile platform, the single board PC and then the PC104 platform was chosen. To allow us to use many of the old IP modules, we needed to develop an IP carrier in the PC104 format. This carrier allows the IP module access to the PC memory via the ISA bus. Since the IPcomm360 has its own flash memory, it is possible to download software necessary to communicate with the lasers. In practice, the telegrams to the lasers can become quite complex and there are subtle timing issues involved with high speed serial communications, in particular, the need to syncronize clocks between sensor and computer. Once the lasers are running at their highest speed, each packet of data transmitted by the laser (single range scan) is copied into the PC memory. It is then time-stamped and an interrupt raised. The advantage

of this approach is that an interrupt is only raised when the entire range packet has been received, rather than whenever the small serial chip fifo is full. This shifts a great deal of communication load from the main CPU to the embedded system. Thus, even with two lasers running at their highest speed, only a few percent of the main CPU is used. Due to the success of this approach, we have plans to build our own 68360 PC104 board, without the need of the IP carrier.

With any data collection it is very useful to have a video record of what the sensors are actually seeing, particularly during data analysis, which may occur several days after the event. For this purpose a video logger was developed. It consists of a DV camcorder, Prolog Time-code generator and video camera. The advantage of this system over a conventional video system is that it is computer controlled and the video tape is time-stamped and syncronized with the time on the computer. This allowed us to leave the system unattended, without the need to keep notes.

3 Software

In the current system the laser data is processed off-line, i.e. data is collected and then processed at a later date. To prevent the software continuously recording data, or requiring manual intervention, a simple threshold is used to detect the presence of a truck below the laser. In practice, a threshold of 8m below the laser was used. When this threshold is satisfied, a series of events are triggered: a file is opened and time-stamped; the video logger is set to record; and range data is logged to the file. Once the truck has passed, the file is closed, the video logger is set to pause, and the software waits for the next truck.

In the field, it is possible to examine the range data from each laser on a laptop (X11 under Linux) as it is being acquired. One laser scans the width of the truck (see Figure 4), the other scans the length (see Figure 5). In both figures, the data is converted to Cartesian coordinates and the laser is positioned at the top-centre of each screen. Both scans were taken when the truck was positioned directly below the laser. Thus, in the width scan we can see the bottom of the tray and the side walls (as in Figure 2). Although it appears that the ground slopes down to the right, it is the laser that is not level. The diagonal lines that extend from the side of the tray to the ground are *laser* shadows, i.e. since the laser cannot see around corners there is an area where the laser's view is obstructed by the side wall of the tray. In the length scan, the headboard and the bottom of the tray are visible. Once again there is a laser shadow extending from the front and rear of the tray to the ground.

The scan of the width of the truck can be used to generate a height-encoded image of the tray. If each scan line from the width scanning laser is filtered and converted to a line of an image file, then an image (see Figure 8) can be generated where the light regions are closer to the laser, and thus, further away from the ground. Although there has been no attempt to

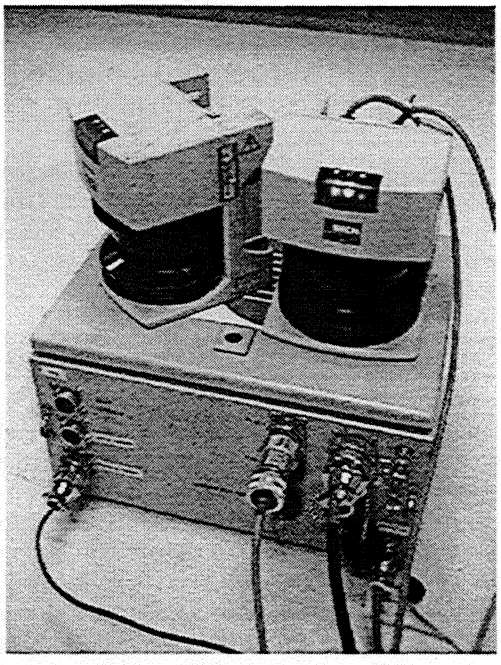

Figure 1: Lasers and Field Computer.

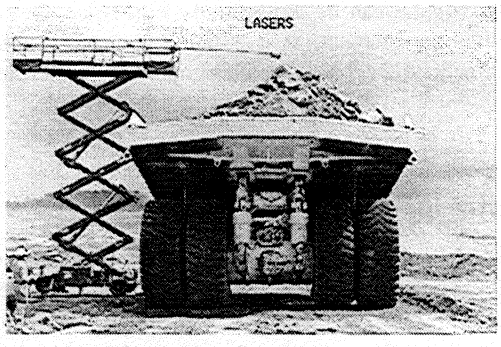

Figure 2: Laser position above truck.

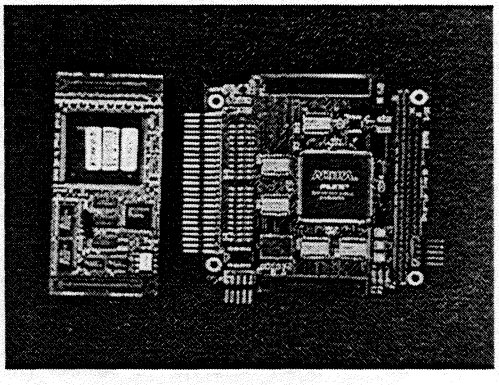

Figure 3: IP module and IP carrier.

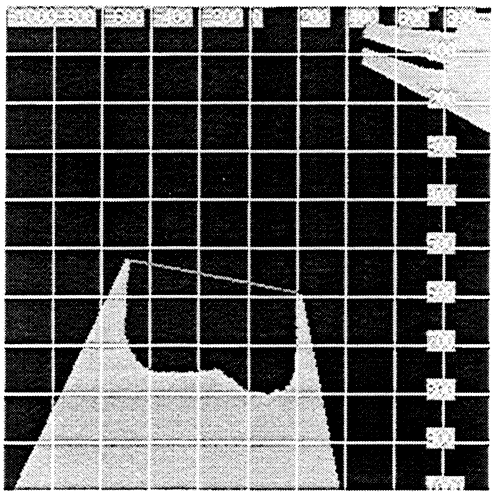

Figure 4: Laser scanning the width of the truck.

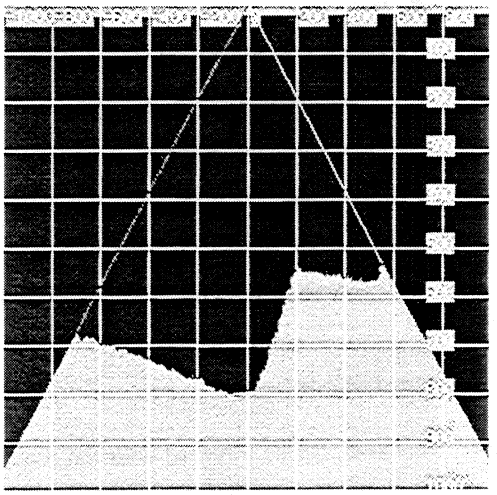

Figure 5: Laser scanning the length of the truck.

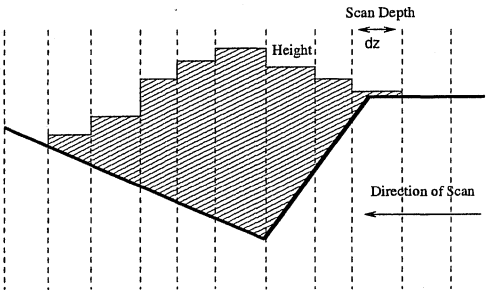

Figure 6: Lengthwise cross section of tray, showing summation of area.

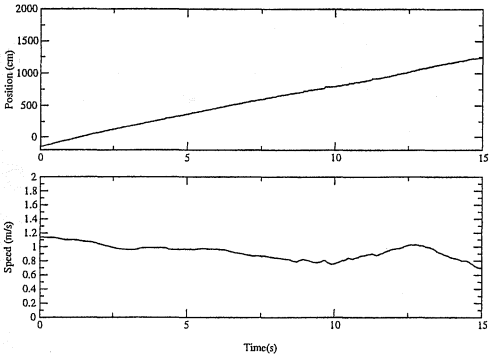

Figure 7: Position and speed of tray.

scale the image, it does show the basic geometry of the tray and features of the truck body (i.e the engine bay). To scale the image, and thus generate a 3D profile, it is necessary to know the distance between each scan (scan depth, dz). This is shown in Figure 6. The distance between scans can be derived from the velocity of the truck. If the velocity of the truck were constant, then each step would simply be the length of the truck divided by the number of scans taken. However the trucks are manually driven and the trucks often brake and accelerate during a scan.

Previously, the velocity of the truck was estimated by following a single feature in the data generated by the laser that scans the length of the tray. The most obvious feature to follow is the front-lip of the headboard (RHS of Figure 5). As the truck moves past the laser, the front-lip of the laser will continue to move to the right (see Figure 7). The speed of the truck is simply: the change in position divided by the time step (0.04s at 25Hz). We are now using techniques based upon Iterative Closest Point (ICP) which are more robust to different tray designs, and occlusions from rocks that can lie on the headboard. By combining the height of the tray with its velocity it is possible to generate a 3D profile (see Figure 9).

The data stored on the data-logging computer consists of several hundred files. Each file represents an event that was triggered by a truck moving under the laser. So the first task was to split the data into those events which were haul trucks and those which were not. In some cases, the laser was triggered by diesel fumes, as the trucks accelerated from the weigh-station. These files were very small and easily removed. In several cases, the trucks were driven too quickly through the laser. Given that the trucks are approximately 12m in length, to maintain a scan step-size of less than 10cm at 25Hz, the truck must remain under the laser for at least 4.8s. At 12m this equates to a speed of less than 2.5m/s or 9km/h. The final task is to identify the different types of tray. This can be done with simple threshold classification, as shown in Figure 10, where three different trays can clearly be identified by the differences in length and width.

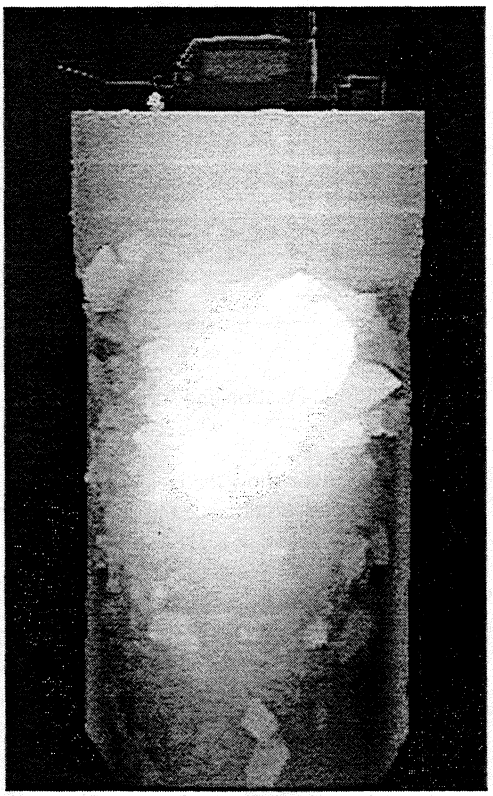

Figure 8: Height-encoded images showing rock fragmentation.

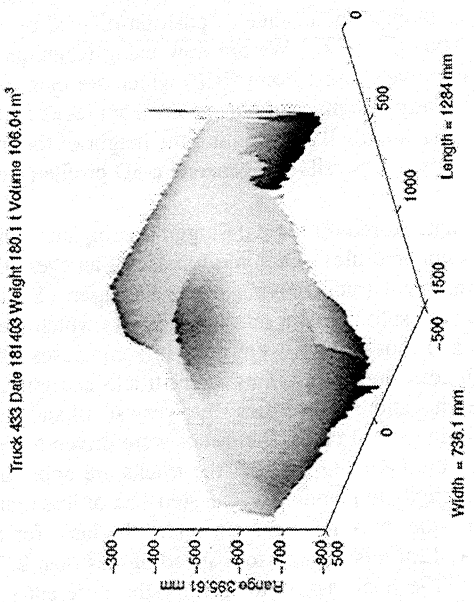

Figure 9: 3D profile of full tray.

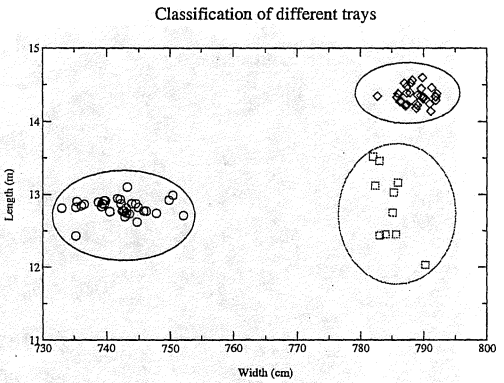

Figure 10: Classification of different trays

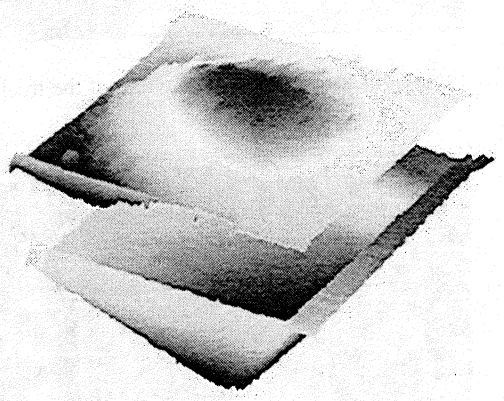

Figure 11: Estimation of volume from 3D profiles.

4 Volume Estimation

Volume estimation can be achieved with the generation of 3D profiles. It can be estimated by subtracting the 3D surface of the empty tray from the 3D surface of the full tray, as shown in Figure 11. Although this is conceptually simple, it is computationally complex. It requires knowing the precise location and orientation of each tray, and a 3D representation of the empty tray. Our software, which is based upon an alternative approach, is faster, more reliable and does not require any calibration. Although the technique cannot be described in detail in this paper (it is currently under a provisional patent) it tracks a few key features. This is shown in Figure 12, where one can clearly see the headboard and bottom of the tray from the range to the centre of the tray. This technique has been successfully applied to 300 trucks, with 7 different trays, on a variety of materials (primarily overburden) at several different mines.

Given the volume of every tray it is possible to characterize different trays with different loads (confirmed with video log). For example, data from one of our field trials indicated:

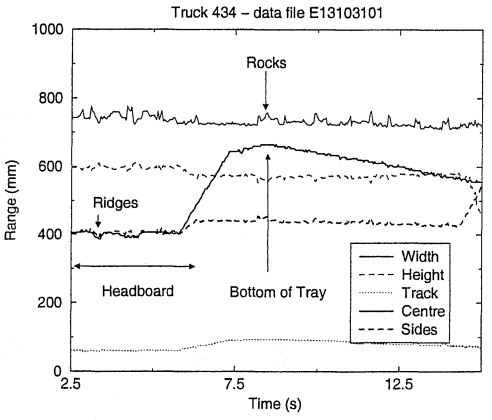

Figure 12: Location of empty tray, from width scan.

Volume	Tray A	Tray B
Shale	95 m^3	130 m^3
Limestone	90 m^3	128 m^3

Here tray B is able to carry 42% more limestone than tray A, but only 36% more shale. This difference is probably due to shovel operation or packing of shale and limestone in each tray. There are many ways such information could be used, e.g. to maximize utilization of this fleet, one should use tray B for Limestone and tray A for Shale. Alternatively, if one were able to match the shovel against the volume in each tray, it may be possible to establish the most efficient combination. Although the primary aim of this technology has been to estimate the volume of material in each tray, there are other features that one can examine and measure:

Load-distribution can be characterised by the centre-of-gravity of each load. Such information can be used for maintenance and shovel design.

Load-profile can be characterised by the slope of the load. Such information can be used to estimate tray utilization and improve design (see Figures 9)

Fragmentation can be characterized by the texture of the 3D surface. Such information can be used to estimate rock breakage for blast optimisation. Two different types of fragmentation are shown in Figure 8.

5 Bulk density

Determining the effective bulk density of any material is not trivial. It is influenced by the density (specific gravity) of the original material; the expansion (swell) when it is removed from the ground; environmental factors (such as moisture) and the distribution of the material in the tray. An effective bulk density can only be estimated if all of these factors are known. Typically, scaling is used to account for the last three

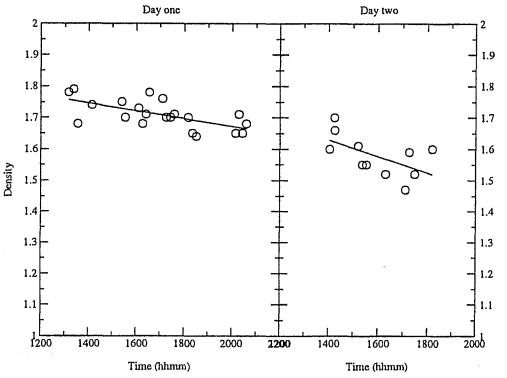

Figure 13: Change in bulk density

factors. In the past, this estimate has been verified by surveying the spoil pile. A more accurate method of determining the bulk density is to simultaneously measure the weight and volume of the material in the tray. This can be done if a weight survey is conducted in parallel to the volume estimation. For example, from a recent trial we determined that

Bulk density	Tray A	Tray B
Shale	**1.79**	**1.80**
Limestone	**1.86**	**1.86**

Another way that the density can be presented is to observe trends in the data over a number of days. This is shown in Figure 13, where the density of overburden falls by 15% over 2 days. Although the material was extracted from the same source it appears that the fall is due to changes in moisture content and fragmentation. Such a fall has significant effects upon truck utilization. For example, if the trays were designed for the first days operation (density of 1.75) then on the second day (density 1.5) they would become under-utilized, i.e. the tray would be full but under weight. Conversely, if the trays were designed for the second day, then on the first day the truck would be carrying too much weight. If this information were available on-line then it would be possible to gradually schedule trucks with larger trays. Perhaps these trucks could be taken from another part of the mine, where the density is going up and smaller trays could be used. It is only with this data that we can fully optimize the truck fleet.

6 Discussion

One of the issues that has become apparent from this work is whether it is necessary to actually know the weight of each load. If the volume is known for tray utilization statistics, then the weight of the load is only required for truck maintenance. In this case, the trucks health can be monitored by stress sensors.

The benefit of this system is the ability to measure the in-situ volume of every haul-truck in the mine. This can be done without the need for the truck to stop. For the first time, ac-

curate production monitoring is possible and productivity can increase on a number of fronts:

- Reschedule trucks: A mine with a variety of trays, trucks and shovels will benefit by rescheduling different haul trucks to different shovels depending on the nature of the material that is being hauled.
- Comparison of tray designs: Trays have been designed to carry material of known bulk density. In a mine these properties change and it is possible to observe how different trays can carry different material.
- Automation of dispatching systems: If both volume and weight is available to the mine-management (in real-time) then changes in truck utilization can trigger changes in dispatching. For example, if a shovel starts to dig material with a high bulk density then trucks with smaller tray should be dispatched, or vice versa.

7 Conclusion

This paper demonstrates the feasibility and advantages of volume estimation. In the current work, the volume of dirt in each tray was calculated off-line, i.e. data was logged and examined several days after the trucks had been scanned. Given the robustness of the system to different trucks, velocities and ore bodies there is no practical reason why this analysis could not be performed on-line. Such a system could be used for an entire fleet of trucks and left in the field as a permanent installation (i.e. on a gantry over the roadway). With additional communications infrastructure it would be possible to transmit the volume to the truck as it is being scanned, or log the volume of each tray directly to some central control where the trucks could be rescheduled to improve utilization.

In the future, we propose to build a new truck scanning platform (see Figure 14) which can be mounted on a mobile gantry. The laser platform will house two lasers, a field computer, a video camera (optional), a radio LAN and a mobile phone. To make the system self contained it will be powered with a solar panel and batteries. The software will be written to automatically detect the presence and speed of each vehicle; distinguish between the various tray designs; acquire image of each vehicle; generate 3D profile and estimate volume of each load. All data will be stored on disk and down-loaded via LAN. If possible, the LAN can also be used to maintain a connection to the mine management system (i.e. directly, or with a web page). The state of the system can be monitored with a mobile phone.

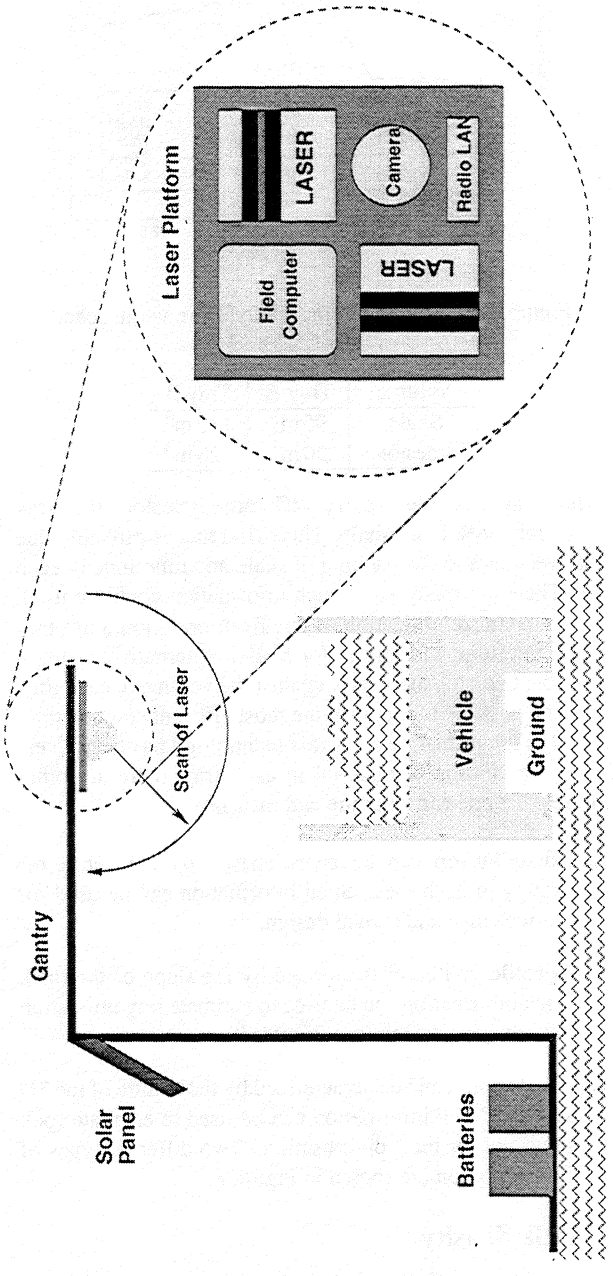

Figure 14: Proposed truck scanning platform.

Acknowledgements

The author would like to thank Pavan Sikka, Leslie Overs, Graeme Winstanley, Stuart Wolfe and Reece McCasker who were either involved in the development of the high-speed laser interface, or assisted in the field trials.