Real-time image processing and data communication for a 3-D measurement system.

D.D.A. Pushpakumara, R.M. Gooch, & T.A. Clarke

Centre for Digital Image Measurement & Analysis,
City University, Northampton Square,
London. EC1V 0HB. U.K.

ABSTRACT

Real-time 3-D measurement based on close-range photogrammetry and machine vision principles has been a key research area at City University for number of years. The focus of this work has been the development of a novel, real-time, high precision, 3-D measurement instrument. Multiple CCD cameras are used to obtain 2-D data concerning the object under measurement to produce 3-D co-ordinates for discrete points on the object. The system is suitable for numerous applications where high precision non-contact measurement is required. A strong emphasis has been placed on ensuring that methods are suitable for use in manufacturing environments. This work covers every aspect of close-range photogrammetry and involves fast electronic data processing and data communications. This paper describes how 2-D image processing and data communication are performed to achieve true real-time 3-D measurement.

1. INTRODUCTION

The process of 3-D measurement using conventional CCD cameras and frame-grabbers is limited due to the image processing and 3-D co-ordinate computational requirements. To circumvent the current bottle-necks, a programme of work has been initiated. This paper considers the 2-D aspects. Complementary projects are addressing high speed correspondence solving and 3-D co-ordinate computation. The title of the complete system has been called *3D-Net* both in acknowledgement of the communications network and photogrammetric network configuration aspects.

3D-Net consists of one or more measurement cells each of which consists of four or more camera stations. Each measurement cell deals with tasks that occur within its limits. The use of overlapping cells, or a distribution of cells centred around each measurement activity enables real-time measurement across multiple manufacturing work areas. The objectives of the *3D-Net* project are as follows:

Real-time 3-D measurement.....Up to 500 targets in 1/25 sec
Minimum latency.....2/25 sec
Flexible architecture..... 3-n cameras per cell, n cells
Low cost.....Hardware cost < £1000/cam excl. cameras
Future proof .....Not tied to a single camera or camera manufacturer
Scaleable ..... Any resolution camera
High accuracy..... Same result as traditional bundle adjustment
Operation..... Simple - Photogrammetrist not required
Easy setting up.....Rule based set up with CAD tools
Automated .....Fully automatic operation from switch on
Redundant .....Failure of cameras is permitted
Robust.....Self checking for blunders on several levels
Continuous operation..... Cameras can be adjusted on-line

In operation the system produces measurement information concerning the 3-D location and orientation of all objects within the measurement environment. The initial specification of tasks is intended to be carried out in a virtual reality environment, monitoring for correct operation is carried out in a CAD model visualisation of the manufacturing process. The operation of the measurement process is illustrated graphically in figure 1.

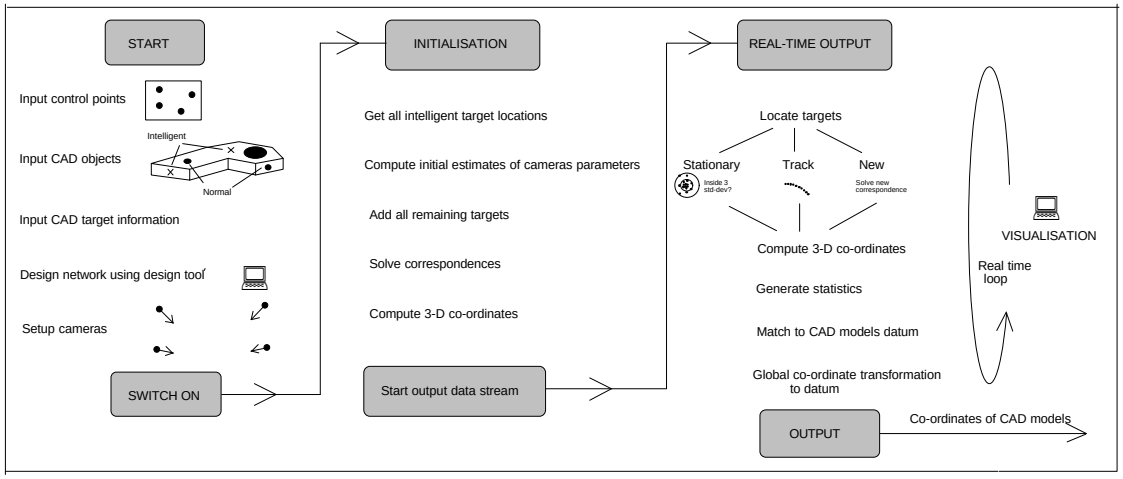

Figure 1. Setting up, initialisation, and real-time operation.

In operation the supervisor can use a 3-D graphics workstation to visualise all operations by selecting the measurement cell of interest and seeing a fully rendered real-time view of the CAD scene in which all the CAD objects are shown in relation to each other from user selected viewpoints.

Efficient acquisition of 2-D data is necessary to achieve the real-time goals. A technique is required to isolate the object of interest from the unnecessary background. In this system, small circular pieces of retro-reflective material (targets) are stuck on the surface of the object (this is acceptable in many applications). The characteristic of a retro-reflective target is that it reflects light back in the same direction as the illumination source. By correctly adjusting the lighting and imaging system, targets appear as bright spots in the image, these can be distinguished from the background using a single threshold.

Real-time 2-D image processing is a major part of the *3D-Net* system. The current prototype system has a single measurement cell and uses five Pulnix TM-6CN CCD cameras of resolution 752 x 568 pixels that operate in interlaced mode. Each camera outputs a video frame every 1/25 of a second. Each camera has a modular processing system attached to it which consists of: a 2-D processing module using a video feature extractor (VFE); a Digital Signal Processor (DSP) based module; an Ethernet module; an GPIO module; and a power supply module. These modules, together with a camera, form an intelligent camera system. The VFE extracts pixel locations of target edges and the intensity. The DSP module uses these data to accurately locate the targets in the image to sub-pixel accuracy. The operation of the DSP module is de-coupled from the VFE by two asynchronous FIFO buffers to achieve parallel operation. The Ethernet module transfers 2-D target location data to the central processing workstations which perform correspondence solving, 3-D co-ordinate computation work, and 3-D visualisation. These tasks require much higher processing power than the 2-D processing tasks due to the complex numerical and graphical processing requirements.

2. VIDEO FEATURE EXTRACTOR (VFE MODULE)

The VFE module is hardware based. It utilises high-speed components and carries out real-time target edge detection, pixel counting, and A/D conversion of the video signal. Figure 2 is a block diagram of the hardware.

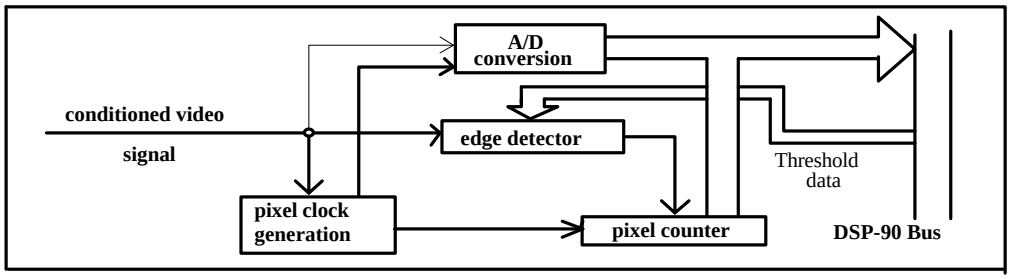

Figure 2. block diagram of video module

The VFE module was developed for use with both line-scan and area-scan video cameras with minor modifications to the pixel clock generation circuitry. A phase-locked-loop (PLL) based pixel clock generator provides the pixel clock signal to the A/D converter and the pixel counter. The 10-bit resolution A/D converter digitises the video signal at pixel rate (every 70nS). A high-speed counter is clocked by the pixel clock to provide the edge location data and is reset at the beginning of each video line. The object edge detector is based on a digital comparator which compares the converted intensity with a pre-selected threshold value to detect the object edge segments. The threshold for the comparator is set by the DSP-90. For each object edge segment, the VFE stores in the FIFO buffers the pixel counter value at the first point that is above the threshold (starting edge) followed by the intensity values up to the point at which intensity falls below the threshold (finishing edge). These are all stored as 16-bit values. Due to the algorithmic requirements described later, data belonging to each video field are stored in separate FIFO buffers. Figure 3 illustrates a portion of an image and the corresponding edge and intensity data that would be produced by the VFE. The objects shown in the image represent a general situation. In practice the retro-reflective targets are both regular in shape and small. However, other high intensity objects in the image must be distinguished from targets. Therefore algorithm considers all possible situations.

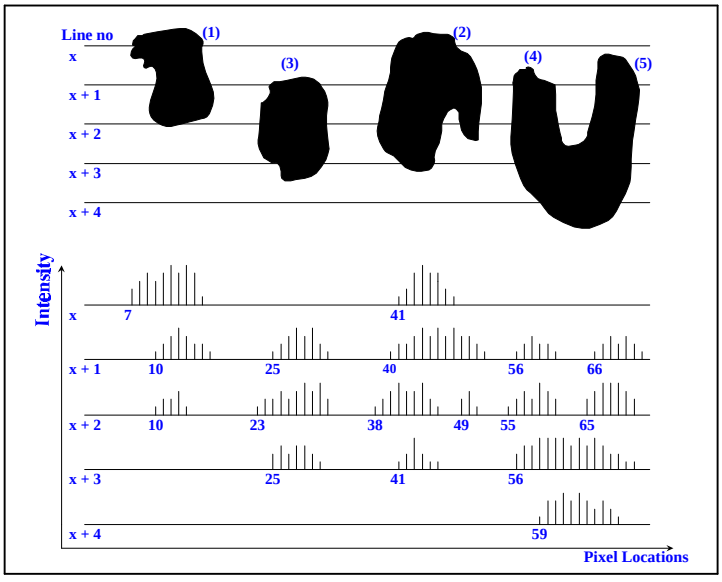

Figure 3. Objects and corresponding edge and intensity data

The beginning of each line of data is identified by a unique number "0" that is placed by the VFE hardware into the appropriate FIFO buffer. One of the high order bits in the data words is used to distinguish between edge and intensity data.

3. DSP-90 PROCESSOR MODULE AND GPIO MODULE

The DSP processor system is comprised of a general purpose I/O (GPIO) board and the DSP processor board itself. The GPIO board has four 8-bit programmable bi-directional I/O ports, one of which is used for setting the threshold of the edge detection comparator. In addition a pair of 16-bit FIFO buffers are provided for storing the edge and intensity data. The processor board uses an Analog Devices ADSP2101 fixed point processor operating at 20 MHz. The ADSP2101 has both its' architecture and instruction set optimised for fast numerical processing. The processor has 16K words (16-bit) of data and 16K words (24-bit) of program memory, of which 2K words of data memory and 1K words of program memory are on-chip. 128 words of the external data memory space is used for memory-mapped I/O ports. The rest of the external data and program memory spaces are filled with 30nS SRAM. The 16K words (24-bit) of boot memory space is occupied by a 70nS EPROM. The processor divides this memory space into 2K words pages (0-7), page 0 contains 2-D processing and data communication software. Upon power up, the ADSP2101 loads a pre-defined page into the on-chip program memory. Software controlled rebooting is possible. In operation the processor reads edge data from FIFO buffers and carries out the object reconstruction, recognition, and location. In addition, the processor controls target location data transmission and data reception via the Ethernet link.

4. OBJECT LOCATION ALGORITHM

The object location algorithm is optimised for use with real-time processes in that it only requires object edge and intensity data belonging to two consecutive lines of the image to carry out processing. It reconstructs, recognises, and calculates the centroid of objects in a progressive manner as each line of data becomes available. The algorithm uses a number of logical tests to ascertain the connectivity of object segments found in the current line with those in the previous line. This requires pixel locations at both edges (starting and finishing) of the object segments. The algorithm derives the location of a finishing edge by adding the number of intensity values to the pixel location at the starting edge as each intensity value represents a pixel location. Two buffers, the current buffer and the previous buffer are used for storage of edge data for current and previous lines. These buffers have 752 locations in each to avoid overflow. The intensity values are directly multiplied with the corresponding pixel values as they are read in. The sum of these products for each object segment are separately stored in a temporary buffer until the connectivity of each segment is established.

The algorithm considers all objects of all possible shapes to ensure reliable operation and isolation of good target candidates. For connectivity determination purposes, the objects encountered can be categorised into five groups: new, finish, continue, split, and merge. If the pixel locations of starting and finishing edges are defined for two consecutive lines (current and previous) as the ***starting_edge_location*** & ***finishing_edge_location*** the logical tests used for connectivity establishment can be described as follows:

1. New object test.

If the previous ***starting_edge_location*** >= the current ***finishing_edge_location***, the object segment under consideration in the current line and belongs to a new object.

2. Finishing object test.

If the previous ***finishing_edge_location*** <= the current ***starting_edge_location***, the object segment under consideration belongs to a finishing object.

3. Continuing object test.

If tests (1) & (2) are not true ***and*** the object segments overlap, they belong to the same object, hence are part of a continuing object.

4. Splitting object test.

If the current ***starting_edge_location*** is < previous ***finishing_edge_location*** the object segment under consideration in the current line belongs to a splitting object. Every continuing object is tested for splitting.

5. Merging object test.

Every continuing object which does not satisfy test (4) is tested for merging. If the current ***finishing_edge_location*** is > the previous ***starting_edge_location*** then the object segment under consideration in the current line belongs to a merging object.

In operation a stack is maintained for object labelling purposes. When a new object is encountered a unique identity number is popped from the stack. The algorithm maintains a pair of buffers for storing the identity numbers of the objects in the current and previous lines which are used for connectivity establishment. In addition another buffer is used in which parameters belonging to the objects are stored. For each new object, a number of locations are allocated in the parameter buffer. These are updated until the object is completely reconstructed (finished). One of the parameters is the sum of the product of pixel locations and intensities, another is the sum of intensities. Other parameters such as: the identity number; the peak intensity; the area; the first line number; and the last line number are stored for target recognition purposes. Parameters for objects that cannot possibly be targets (objects that are large or of irregular shape) are not updated. However, each object is completely traced to retain the integrity of the recognition process but any parameters pertaining to these objects are discarded when the object is completed. When an object passes the criteria for selection as a target the identity stack is pushed and grey-scale centroid is calculated and any other pertinent information for future target related purposes are stored. The locations in the parameter buffer which are then freed are then made available for a new object.

The buffers and stacks require 9K words (16-bit) of RAM and are implemented in the external data memory of the DSP-90 system. An executable version of the algorithm coded in the ADSP2101 assembly language takes 892 words (24-bit) of program memory. The processor takes approximately 67μS per target to reconstruct, recognise, and compute the centroid of a target which is assumed to be 5 x 5 pixels in size.

5. ETHERNET COMMUNICATION LINK

The IEEE 802.3 compliant 10Base-2 Ethernet module is based on the National Semiconductors DP83901 Network Interface Controller (NIC). Manchester encoding is integrated in the NIC. The transceiver is developed as a separate module to suit other applications. Figure 5.1 illustrates the block diagram of the hardware architecture.

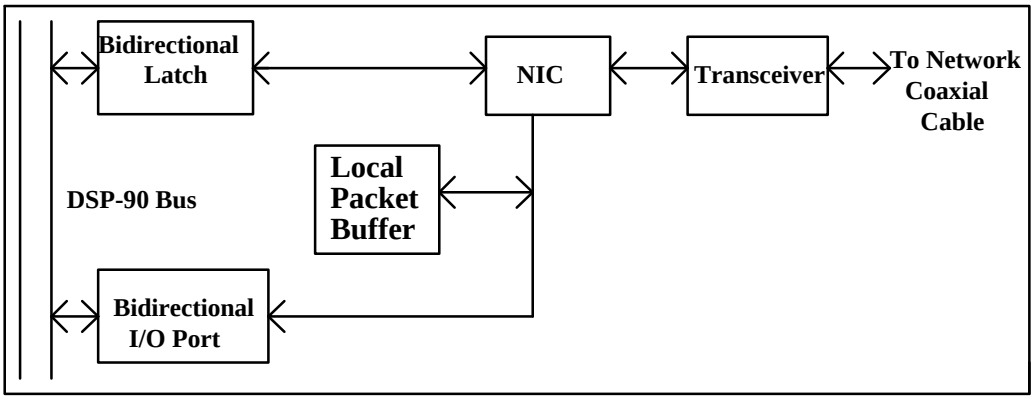

Figure 4 Block diagram of Ethernet module

The activities of the high priority network are decoupled from the DSP-90 bus using a 16-bit bi-directional I/O port. The I/O port based architecture used with a local packet buffer avoids the maximum bus latency requirements of the NIC. Packet transfers to and from the module are via this port. The module has a 16K x 16 bit local packet buffer RAM which is divided in to two sections for receive and transmit operations. When a packet is received, the NIC's local direct memory access (DMA) channel transfers it into the receive section. The DSP-90 is informed of the reception of a packet by means of an interrupt that is generated by the NIC. Consequently, the DSP module informs the NIC of its' intention to read the received packet into data memory. The NIC then reads the packet, a word at a time and writes (remote DMA transfer) to the bi-directional I/O port from which DSP reads. The NIC provides the necessary hand-shake signals for this purpose. The transmit operation is relatively simple. The packet is placed in the local packet buffer using the remote DMA facility of the NIC. It is then transferred to the network via the local DMA channel. The bi-directional latch is used for reading and writing a 8-bit register data to on-chip NIC registers. Figure 5 illustrates the data communication network.

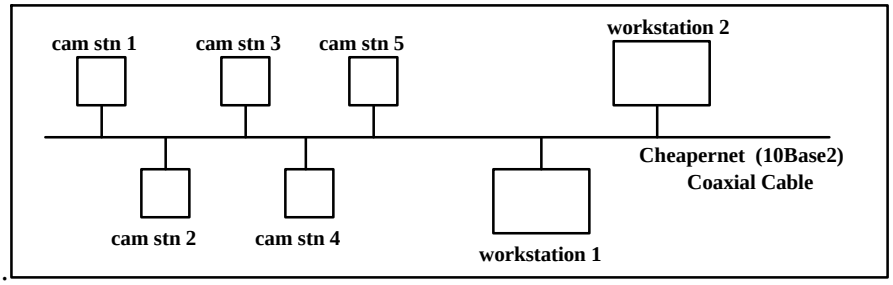

Figure 5. Data communication network

Low level software has been developed for packet transfer and receive operations. Assembly coded routines consume 0.7mS of CPU time for either operation. An efficient protocol for real-time data transfer between camera stations and the central workstation is under development. The basic operation of the protocol is that the workstation broadcasts a message to all camera

stations asking them to output their latest target location information for the most up-to-date image that has been processed. The camera stations which continually process the image data, also monitor the transfer instruction. Acknowledgements are then sent and the target location data is transferred to the workstation. A possible problem here is that multiple camera stations may compete to gain access to the transmission medium resulting in excessive collisions. As the number of targets and camera stations increases, the probability of failure rises. However, the data rates involved are relatively low. Even with each camera viewing as many as 300 targets, the target location data (X,Y co-ordinates, 16-bit each) for a single video frame are unlikely to exceed the IEEE 802.3 Ethernet standard packet size. The protocol uses an ID number based priority scheme. It allocates a priority ID number for each camera so that the camera station with highest priority would transmit first. Priority allocation could be based on many factors such as camera location and number of targets that a particular camera sees. Work on this aspect is still at an early stage.

7. INTELLIGENT CAMERAS

To obtain high accuracy 3-D measurements simultaneously at several locations in a large scale manufacturing situation requires that many groups of cameras are used to obtain convergent views of the objects to be measured. This means that the issues of reliablity, robustness, and redundancy are paramount. *3D-Net* has been designed around the concept of intelligent cameras using the hardware and software described in this paper. This requires that each camera can be added to the network or removed from it without interrupting operations. Furthermore, intelligent cameras must also be able to perform low and high level tasks - for instance target tracking assists greatly in lowering the correspondence problem overhead, and correcting image co-ordinates using prior calibration data cuts down the 3-D computational load. Figure 5 illustrates the electronics required for an intelligent camera station without the housing. The electronics in each compact PCB module have a 90 mm. diameter circular format with snap-off ears to enhance mounting options. A common DSP-90 bus runs along the stack-through connectors. The Ethernet and VFE modules use surface mount technology to achieve the small footprint of the complete design. A switched-mode power supply module provides power to the complete system from a wide-band DC input.

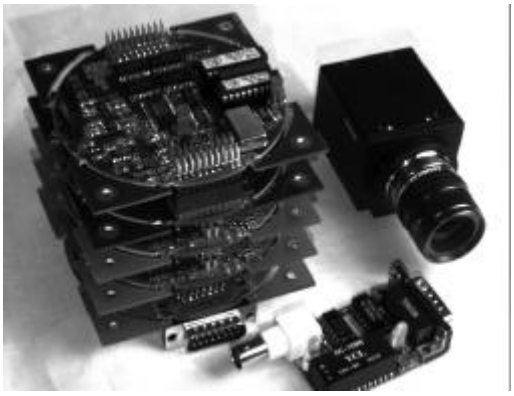

Figure 6. Intelligent camera

It is not intended that the camera and processor should be physically combined as this would limit the design to the camera chosen, rather the processor will be boxed and mounted close to the camera. The reason for this is to maintain the maximum flexibility such that any CCIR camera or pixel clocked camera can be used with the system. Modifications to interface to a range of digital cameras are relatively trivial and will be made in the near future.

8. CONCLUSION

Using off-the-shelf low cost components a novel system has been realised. The total cost of this system is less than £1000 (excluding camera). The basic hardware can be used in numerous other applications such as robotics, and visual inspection. The system will be used in an underground sewer inspection profiler, and in a geotechnical centrifuge. With respect to high precision 3-D measurements, the system can be further developed. The relatively low resolution CCD cameras can be replaced with high resolution digital cameras to achieve enhanced precision with the minimum of hardware changes. In fact, half of the VFE module would not be not necessary in this case. The DSP-90 system will be able to deal with these larger images due to the lower frame rates which are typical. Further, the use of non-interlaced digital cameras will enable even more efficient implementation of real-time thresholding of intensity levels. This paper has described some of the elements of a system which is believed to deliver a better cost/performance ratio than many comparable systems as well as the significant advantages of a distributed system with parallel computation.

9. ACKNOWLEDGEMENTS

The author would like to gratefully thank: National Semiconductor for providing Ethernet components and literature, and the Worshipful Company of Scientific Instrument Makers who are currently sponsoring the PhD of D.D.A. Pushpakumara.

10. REFERENCES

1. Gooch R.M., Clarke T.A., Pushpakumara D.D.A., & Ellis T.J. '*Driving a DSP underground'* Proc DSP'96, Denmark 1996.

2. Shortis M.R., Clarke T.A., & Short T. '*A comparision of some techniques for subpixel location of discrete target images'* Proc SPIE **Vol. 2350** Videometrics III,Boston, Massachusetts, November 1994, PP 239-249.

3. Clarke T.A. *'An analysis of the properties of targets used in digital close range photogrammetric measurement'* Proc SPIE **Vol. 2350** Vodeometrics III, Boston, Massachusetts, November 1994, PP 251-260.

4. Ellis T.J. *'A programmable processor for real-time feature extraction'* Proc SPIE **Vol. 848** Intelligent Robots and Computer Vision VI, Cambridge, Massachusetts, November 1987, PP-594-599.

5. Pushpakumara D.D.A. *'Msc Thesis'* City University, London 1995.

PAPER REFERENCE

D.D.A. Pushpakumara, R.M Gooch, & T.A. Clarke, 1996, Real time image processing and data communication for a 3-D measurement system. 15th National Information Technology Conference, Computer Society of Sri Lanka. pp. 1-9.