3D imaging and ranging by time-correlated single photon counting

by A. M. Wallace, G. S. Buller and A. C. Walker

3D imaging is an important tool for metrology and reverse engineering of components and structures in a wide variety of contexts, from automobile and aeroplane manufacture to the creative media and architectural surveying. In this article, we review briefly the principal methods in current use for 3D imaging, then present a new method for time of flight depth measurement, which is more accurate and sensitive than the current techniques. To illustrate its potential, we show a number of examples of 3D data acquired from both small and large objects, taking examples from cars, planes and archaeological artefacts.

There is an increasing need for 3D imaging systems to acquire accurate range and image data for a variety of industrial applications. These include accurate surface mapping for metrology and reverse engineering in the automotive and aerospace industries, scanning of large structures for architectural or archaeological surveying, and the creation of virtual and augmented reality environments for leisure, education and simulation. For a number of years we have been working on the design of triangulation and time-of-flight laser sensors for depth imaging, and the subsequent processing of that data for visualisation, segmentation, object recognition and pose definition. Much of our work has been in collaboration with industry; in particular we have enjoyed very profitable associations with British Aerospace, 3D Scanners Ltd. and Edinburgh Instruments Ltd. In this article we discuss briefly the principal characteristics of current approaches to depth imaging based on triangulation and time-of-flight, then present an improved approach based on time-correlated single photon counting.

Passive and active systems

Broadly speaking, optical 3D imaging systems can be split into two categories. Passive systems rely on the acquisition and processing of two or more intensity images acquired from different viewpoints. Provided that corresponding points or features can be identified within the two or more images then it is possible to use the laws of projective (or other) geometry to recover 3D information from the scene. The correspondence problem is not trivial. To achieve high accuracy, photogrammetric measurement systems rely on the placing of calibrated targets on the surfaces in question which can be clearly identified in each image. Where this is not possible, for example in aerial imaging, window correlation or other techniques can be used to identify the corresponding points.

Once corresponding points are identified, the difficulty of the 3D mapping process depends on how much is known about the calibration of the optical system and the structure of the scene. For a calibrated system in which all the intrinsic camera and extrinsic placement parameters are known 3D mapping is a simple process, but researchers in computer vision have made significant progress in recent years to recover 3D data using uncalibrated systems.[1] Apart from the theoretical interest, this has important uses for 3D navigation by mobile robots (whose position is not known) or 3D reconstruction from archival photographs (where little may be known), to name but two examples.

Active 3D sensing systems use the projection and scanning of an energy source, usually a laser, onto the scene, then acquire one or more images of the scene subject to the laser illumination. There are many

variations on this theme,[2] but to get a full 3D image it is necessary to project the laser over a wide area. This can be accomplished either by projection of a 2D pattern/texture onto the scene, scanning a stripe across in the scene in one dimension, or scanning a spot across the scene in two dimensions. Image capture is naturally synchronised to the scanning process. A wide range of active 3D imaging techniques are applied in industry, of which interferometry and fringe projection systems are often favoured because of the high accuracy that can be achieved from co-operative targets. However, they are very labour intensive and require the use of mirror reflectors at worst or uniformly reflecting surfaces at best to attain acceptable performance. As a result there has been a search for fully automated 3D scanning systems that can acquire dense and accurate data directly from a wide range of scenes and surface materials. The main competing technologies are laser triangulation and time-of-flight laser systems.

Triangulation

Fig. 1 illustrates a triangulation system that we have constructed at Heriot-Watt University. The source laser beam (not visible in the picture) is directed through a series of lenses, to focus and stretch the beam in one dimension, so that a plane of laser light is projected vertically onto the object surface, in this case a stepped calibration block. The intersection of the laser plane with the object surface produces a stripe in the image of a camera. As the optical axis of the camera is at an angle to the laser plane the horizontal displacement of the stripe in the laser image is proportional to the distance from the object. This can be seen clearly in the displayed image taken from the camera in the left hand of the researcher. To produce a full 3D image, either the laser plane or the object must be translated. In the system shown, the block is sitting on a linear translation stage, which is controlled by the same computer, which acquires the image data, so that data capture is synchronised to the lateral motion

Fig. 1 Conventional flat-bed triangulation scanner. The displacement of the stripe in the image is proportional to the distance to the object

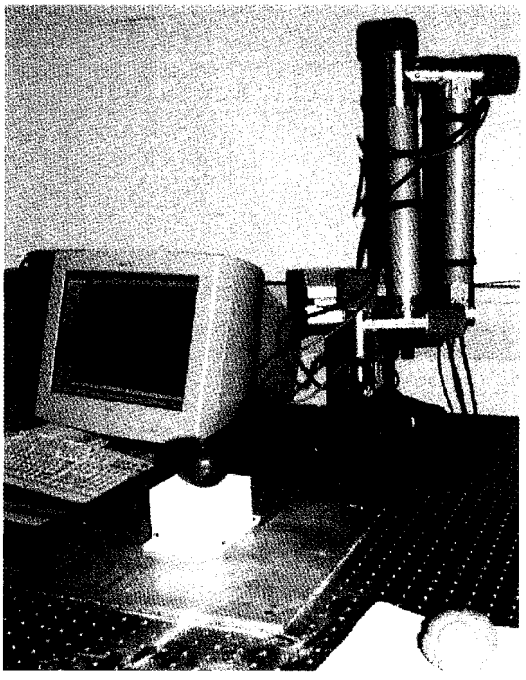

Fig. 2 The 3D scanner's triangulation head is mounted on a Faro arm. This allows measurements to be made all round the object in a larger working volume while maintaining accuracy

and a 3D image is acquired automatically.

Triangulation is a mature technique, but there are several outstanding problems. First, there is an inevitable trade-off between the scale of the object measured, the stand-off distance and the accuracy of measurement. Second, there can be occlusion at concavities due to the necessary separation of viewpoint between the viewing camera and the laser projector. Further problems can occur with poor and false returns due to variable material reflectance and multiple reflections.

Fig. 2 shows an example of an existing commercial system made by one of our collaborators, 3D Scanners Ltd. This uses a co-ordinate measuring arm, from Faro Ltd., and has a triangulation scanning head mounted at the end of the arm in close proximity to the usual touch probe. This is a useful engineered compromise to the geometrical problems discussed above. It is possible to move the triangulation head to a number of positions around the object to be scanned, and use the known kinematics of the robot arm to determine the position and orientation of the head in 3D space. Hence a complete, 'all-round' point cloud of 3D data can be acquired. Typically, this maintains an accuracy of approximately 0·1mm over a working volume of $0{\cdot}5 \times 0{\cdot}5 \times 0{\cdot}5$m. Fig. 3 shows an example of data acquired from this system. The bust is a depiction of James Watt; although only a frontal view is shown, this is a full, shaded 3D mesh constructed from a dense cloud of 3D points.

Time of flight

In contrast to triangulation, a time-of-flight (TOF) ranging system acquires absolute 3D measurements along the line of sight of a common transmitter-receiver optical axis. This immediately eliminates the problems of occlusion and of fixed baseline geometry, which can cause problems with the former technique. Unlike interferometric systems that also act along a single line of sight, the goal is to measure directly from an uncooperative surface without recourse to the use of targets or manual intervention. In order to measure time elapsed, the laser beam is modulated; there are a number of research prototypes and commercially available TOF systems that rely on either amplitude, frequency or pulsed modulation.[3-5]

Fig. 4 illustrates the basic principle of a pulsed TOF system. A laser signal is projected towards the object surface and the reflection is acquired by a receiving optical system. As the speed of light in the transmitting medium (usually air) is known then the distance to the surface can be measured accurately as $\frac{1}{2}c\tau$ where c is the speed of light and τ the elapsed time.

The lower part of the Figure shows a common operating principle. The laser pulse is directed towards the target and a proportion of the signal is diverted to a photodiode or other detector. This start signal is amplified and used to trigger a ramp, which is the basis of the time interval measurement. The returning signal from the target is detected by the same or another photodiode and conditioned to stop the ramp. The signal amplitude provides a measurement of the time elapsed and hence distance to and from the target. This level can be digitised at appropriate resolution and successive values can be stored as the beam is scanned over the surface to build up a complete 3D image.

As an alternative to timing by the analogue ramp a digital counter can be used, but this

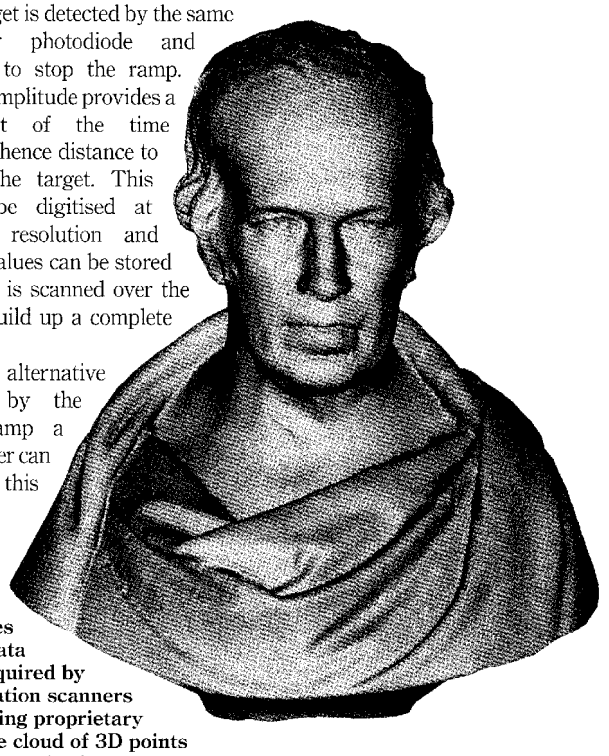

Fig. 3 A bust of James Watt. The data has been acquired by the triangulation scanners of Fig. 2. Using proprietary software, the cloud of 3D points has been polygonised to create a 3D mesh, smoothed then rendered as a shaded 3D model

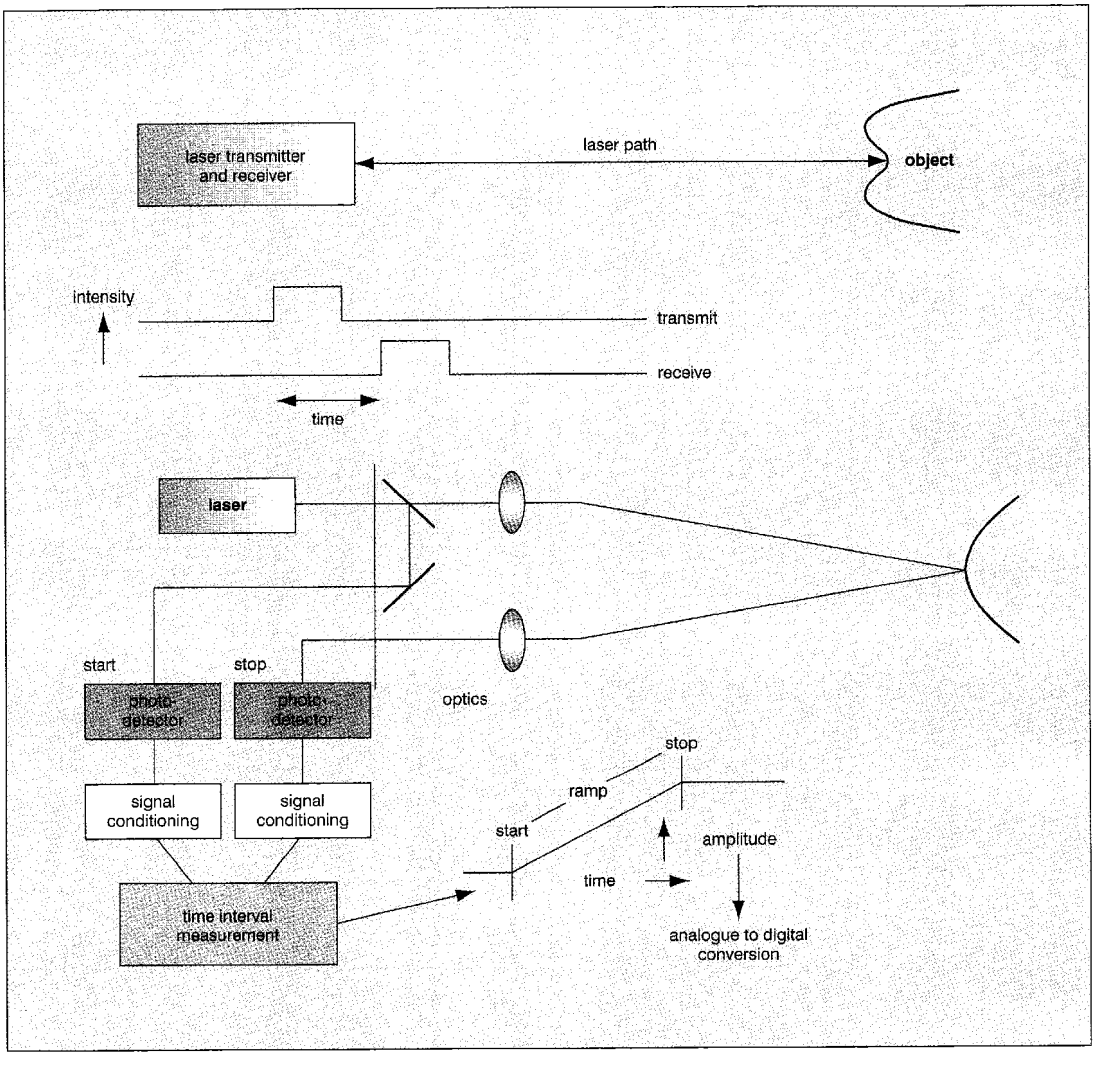

Fig. 4 Schematic diagram of a pulsed time of flight system

is generally less accurate. An amplitude modulated TOF system is conceptually similar. The amplitude of the beam is usually modulated sinusoidally at a fixed frequency so that a measurement of phase shift between the transmitted and received signal provides the equivalent measure of elapsed time. The majority of FM systems have used the self-mixing effect, in which interference between an internal (to the sensor) cavity and an external (between sensor and object) cavity causes modulation of the output power that is a function of the modulation frequency and distance to target. If the frequency is modulated in a controlled manner, then the range can be computed from the observed output power fluctuations.

There are a number of important factors in maintaining the timing resolution and hence distance accuracy in TOF ranging systems. These include the shape (rise time, full width half maximum) of the laser pulse, and the efficiency and response time of the photodetector. For signal conditioning and measurement, the linearity, or at least calibrated repeatability of the ramp, minimisation of timing jitter and the resolution of the ADC and digital timing circuit must all be optimised.

Assuming that these design considerations can be satisfied, existing TOF ranging systems are accurate to about 0·1mm in depth, depending on how the measurement is made. However, this accuracy is, like the triangulation systems, dependent on receiving a relatively strong signal from the target that can be detected and processed. This is further complicated by the problem of range-intensity crosstalk, namely that in all of the modulated signals the time of detection of the returned signal is influenced by its amplitude and hence the reflectance properties and range of the target. A

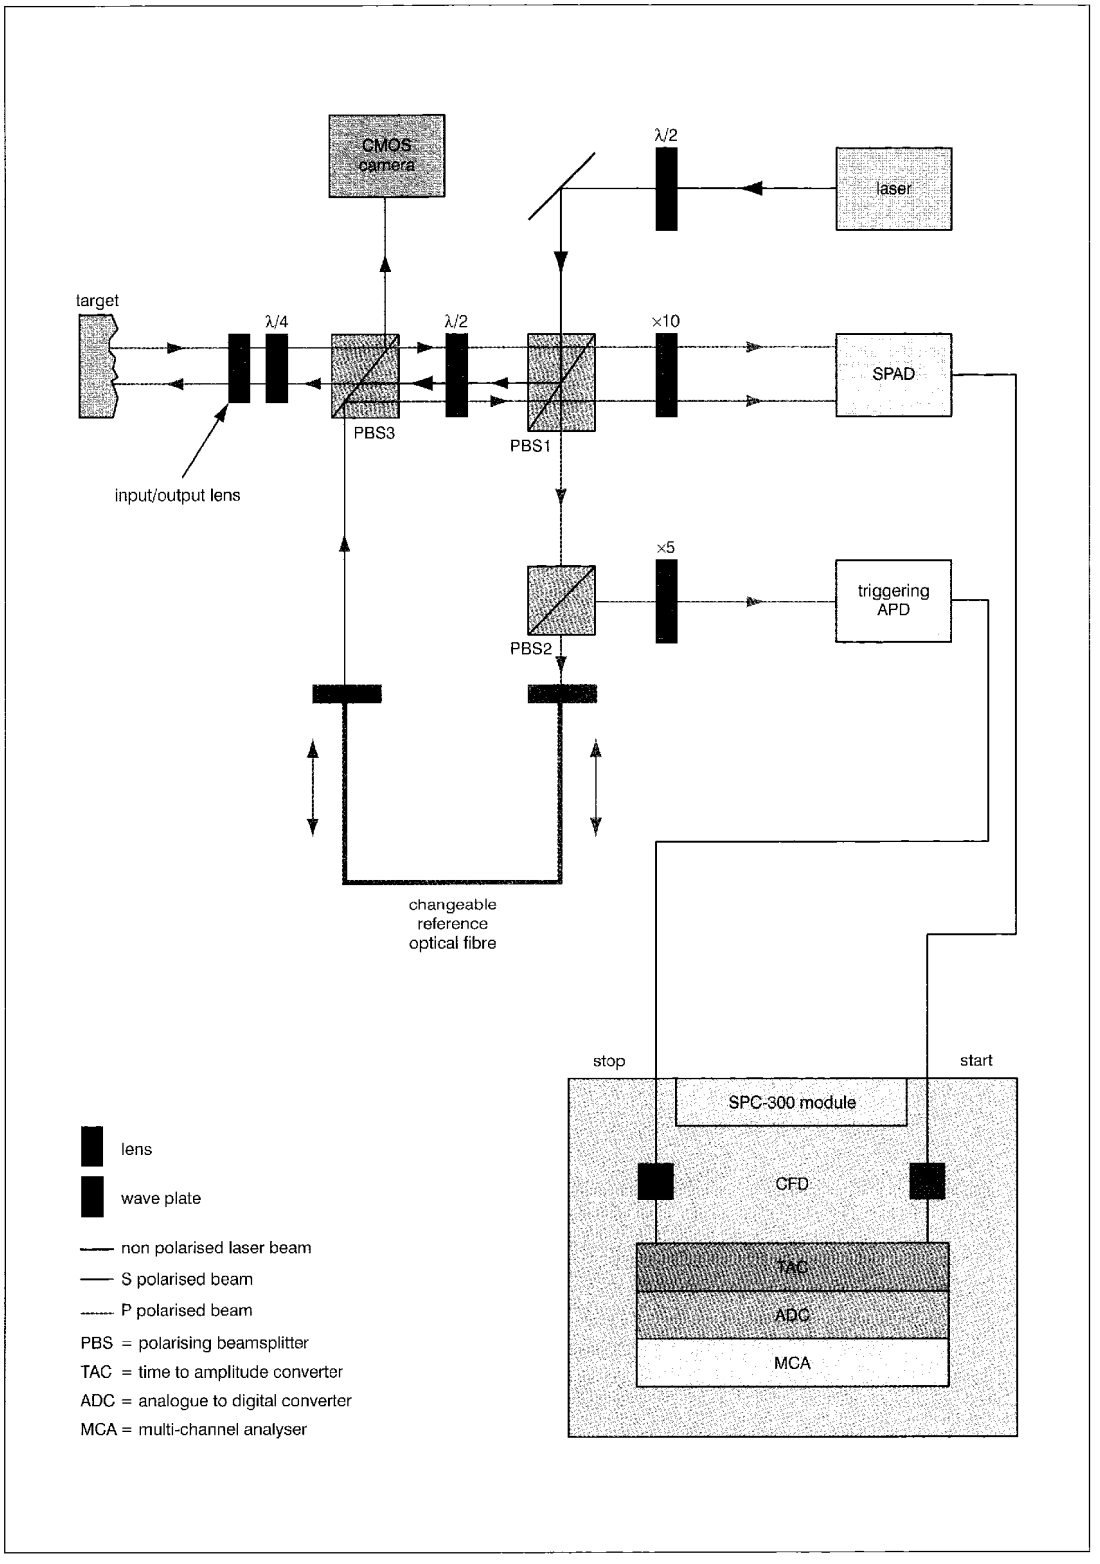

Fig. 5 Schematic diagram of the TOF-TCSPC system

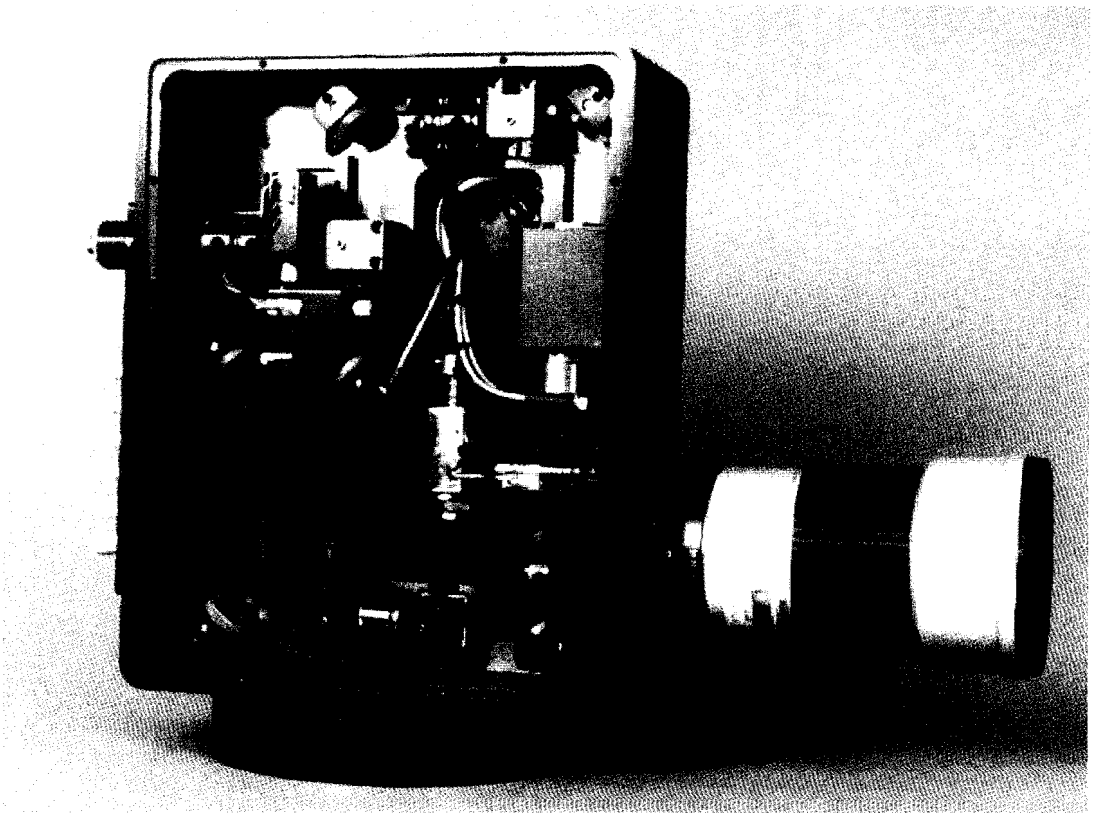

Fig. 6 Photograph of the TOF-TCSPC sensor. This shows the layout of the optical components on one side of the optical baseplate

further, often unacknowledged, problem is temperature sensitivity, although this can be alleviated either by temperature control or by use of a reference channel for calibration. These limitations have led us to develop a new approach to TOF imaging based on time-correlated single photon counting.

Time-correlated single photon counting

Time-correlated single photon counting (TCSPC) is a statistical sampling technique with single photon detection sensitivity, capable of sub-picosecond timing resolution. This technique offers several advantages when compared with previous methods for 3D ranging and imaging based on time-of-flight, i.e. very accurate time (and hence distance) resolution, very high sensitivity, and the ability to cope with variation in amplitude of the reflected return of several orders of magnitude.

Fig. 5 shows a schematic diagram and Fig. 6 a photograph of the TOF-TCSPC sensor.[6] The laser source is a passively Q-switched AlGaAs laser diode (developed at the Ioffe Institute, St. Petersburg) which emits 10-20ps pulses of energy ~10pJ at 850nm and a repetition frequency of up to 25MHz. The major part of this laser signal is directed toward the target by the polarising beam splitter, PBS1, and the back-scattered light is subsequently captured by the input/output lens, in this case a conventional telephoto lens for a single-lens reflex camera. At PBS2, the remainder of the laser signal is split again. A fraction of the laser pulse is diverted to a reference channel, in this case one optical fibre from a set with appropriate lengths (delays).

At PBS3, the received signals from the target and that passed through the optical fibre are combined and transmitted back through the series of waveplates and beamsplitters to be imaged onto the actively-quenched single photon avalanche diode (SPAD). Meanwhile, the final fraction of the source signal is sent to an avalanche photodiode (APD), which acts as a trigger to a single photon counting (SPC-300) electronics module. The temporal location of this reference pulse provides the START input for a time-to-amplitude converter (TAC). The signal from the SPAD provides the STOP input to the TAC, and it is these returns, from both the target and the fibre-optic delay (reference), that are used to build up a photon detection histogram of the type shown in Fig. 7. The final component in Fig. 5 is a conventional CMOS imaging camera. This provides an intensity image of the

target in the region surrounding the laser spot and is useful for visual monitoring of alignment, particularly as the laser (850nm wavelength) is not visible to the human eye.

The SPAD is a key element of our system; the response consists of a fast rise and a slow wavelength-independent exponential tail. We have used devices developed at the Politecnico di Milano, which have a particularly high specification, although commercial alternatives are now available from a number of suppliers. Measurement accuracy is affected by statistical fluctuations in the avalanche build-up time, which cause peak width variations, the maximum electric field in the active junction, which affects the time resolution, and thermally generated carriers in the depletion layer, which result in an inherent dark count. The tail is due to minority carriers, photogenerated in the neutral region beneath the junction, that succeed in reaching the depletion layer by diffusion.

Fig. 7 represents the number of detected photons (vertical axis) against time (horizontal axis). The horizontal axis is defined as sampled channels; in this example each channel corresponds to 2·44ps. The two peaks correspond to the accumulated target (left) and reference (right) single photon returns. It is important to emphasise that each recording on the vertical axis of Fig. 7 represents a single photon event, i.e. the first received photon from either the target or the reference channel in response to a burst of laser energy in the form of a pulse. The average count rate is intentionally much less than one per pulse, ensuring equal probability of recording a photon arriving at any time during the interval. The number of photons collected depends on the laser repetition rate, the acquisition time and the properties of the optical system and target surface.

Using the difference between the target and reference signals recorded by the SPAD to define the distance, rather than using the APD, enables us to balance any temporal variation in the SPAD signal channel. Rather than simply measuring the peak positions we process the histogram data to obtain an accurate estimate of the time lag between the reflected target and reference signals in the histogram. As with the previously described pulsed system, the distance between the target and reference is $\frac{1}{2}c\tau$, where c is the speed of the laser light and τ is the time lag.

Processing the TCSPC data

The number of counts in each channel c_i of Fig. 7 follows a Poisson probability distribution with mean and variance μ_i, where μ_i is the number of observed counts in channel c_i. The pulse histogram is contaminated with a few spurious (dark and stray photon) counts, but in practice the number of ambient light photons picked up by the system is very small. This is because the camera has inherent discrimination: spatially (through lens focusing), temporally (through the use of a finite window

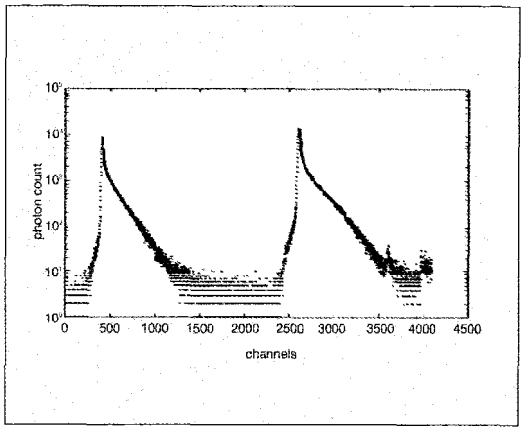

Fig. 7 Histogram of the accumulated photon count (vertical axis) against time (measured in 2·44ps channels, horizontal axis)

initiated by the APD), and spectrally (through the use of interference filters). Another potential source of distortion of the observed pulse histogram is 'pulse pile-up'. This can occur because, for each laser pulse, there is a finite probability that more than one photon will arrive at the detector during the time window. In that case, the first received photon is registered, causing a 'pile-up' in the lower (early) channels. To eliminate this effect we first operate typically at 5% or less photon return rate, so that the chance of more than one return is small, and second we apply a statistical correction factor to allow for the still finite probability of more than one photon.

The histogram can be considered as a probabilistic density function (PDF). In practice we can work with either the raw histogram or an auto-correlation function. In the former case it is the separation of peaks, and in the latter case the position of the sidelobe that gives the elapsed time and distance to the target. We have

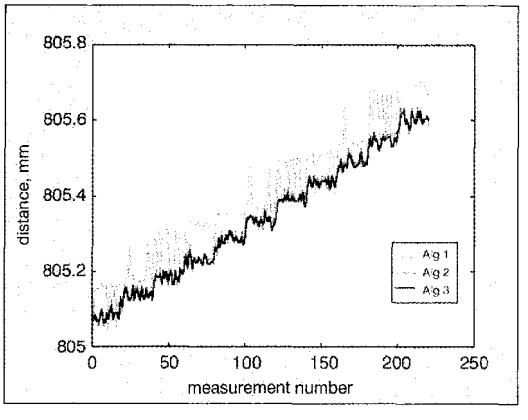

Fig. 8 Comparison of depth data measurements on 50μm steps in the measurement plane. Alg1: centroid comparison. Alg2: upsampling followed by centroid comparison. Alg 3; fitting of operating model to autocorrelation function

experimented with a number of methods to determine this time, working either with raw histogram or auto-correlation data, and using sub-channel reconstruction to improve accuracy.[7] This work is on-going, but we have found that the following procedure, usually applied to the auto-correlation data, has been effective in processing well-defined peaks. As the exact functional form of the histogram is unknown, a nearest representation or 'operating model' has been used to describe the underlying true distribution. This has the form:

$$f(c_i) = \left[h(c_i,\gamma)\left\{1+\sum_{k=1}^{m} a_k p_k(c_i)\right\}\right]/\beta$$

$h(c_i,\gamma)$ is a symmetrical Lorentzian function of width γ. Since the sidelobe is in general assymmetrical this key function is multiplied by a polynomial function, selected from the first m Hermite or Laguerre polynomials. The product of these produces a model that is general enough to fit the empirically observed form of the received data. The best key and polynomial function parameters are determined by a maximum likelihood estimate (MLE). The iterative fitting algorithm employed is as follows:

Repeat

- Fit the parametric key function to the sidelobe of the autocorrelation function of the target and reference single photon returns histogram.
- Estimate the polynomial coefficients which are conditional on the estimated parametric key function parameters.
- Estimate the parametric key function parameters which are conditional on the estimated polynomial coefficients.

Until (convergence)

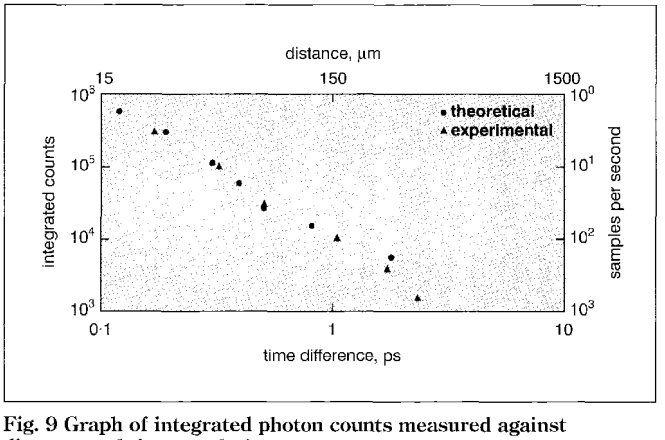

Fig. 9 Graph of integrated photon counts measured against distance and time resolution

Evaluation and application of the system

We have conducted extensive evaluation of the performance of the TCSPC-TOF system with regard to longitudinal (z) and transverse (x,y) spatial resolution, single point measurement time and long-term stability. The longitudinal spatial resolution is determined by the accuracy of the timing and data processing described in the previous sections. In one experiment, we employed a metal target mounted on a micropositioner at a distance of just over 800mm from the sensor. The resolution of the micropositioner was 10μm; we moved the metal target in steps of 20μm toward the sensor and verified the distance moved with a touch probe (Mercer 122). At each position we took a set of 20 measurements with the TCSPC-TOF sensor. The whole experiment was repeated 11 times at 20μm step size, then the same procedure followed for step sizes of 30, 40 and 50μm. The mean of each set was used as the true measurement of the sensor against the ground truth determined by the touch probe.

Fig. 8 shows the data for the 50μm step size data. The three sets of data (Alg 1,2,3) refer to the use of three different approaches to interpret the histogram. The first method used the computed centroid of the sidelobe as a measure of position on the time axis; the second used up-sampling to improve the centroid estimate. The Figure shows the superior accuracy of the fitting of the operating model (Alg 3) described above in comparison with the simpler forms and demonstrates a longitudinal accuracy of 15μm.

Transverse resolution is determined by the spot size; for the current set-up the $1/e$ spot size is 60×170μm at a distance of 2m and ~400μm at a distance of 13m. For a diffraction-limited system the expected sizes would be 60μm and 380μm, respectively. Of course the 3D measurement accuracy is limited also by the precision of the scanning system; we use a pan-and-tilt head supplied by British Aerospace which is accurate to 5 arc seconds. This is comparable to the spot size, for example 48μm at 2m. As regards measurement time, we have found that there is a trade-off between the accuracy and time of measurement, illustrated in Fig. 9. This shows the accuracy of measurement as 15μm with about 10^6 integrated photons, which corresponds to ~1s measurement time, worsening to 150μm at 10^4 integrated photons, which could be accumulated in 0·01s. Improving this figure is a subject of current concern as we are experimenting with improved methods to process low photon count rates to get more accurate measurement.

Long-term stability is governed primarily by temperature variation; like most metrology systems this can cause problems so we have correlated our distance measurement against temperature in the

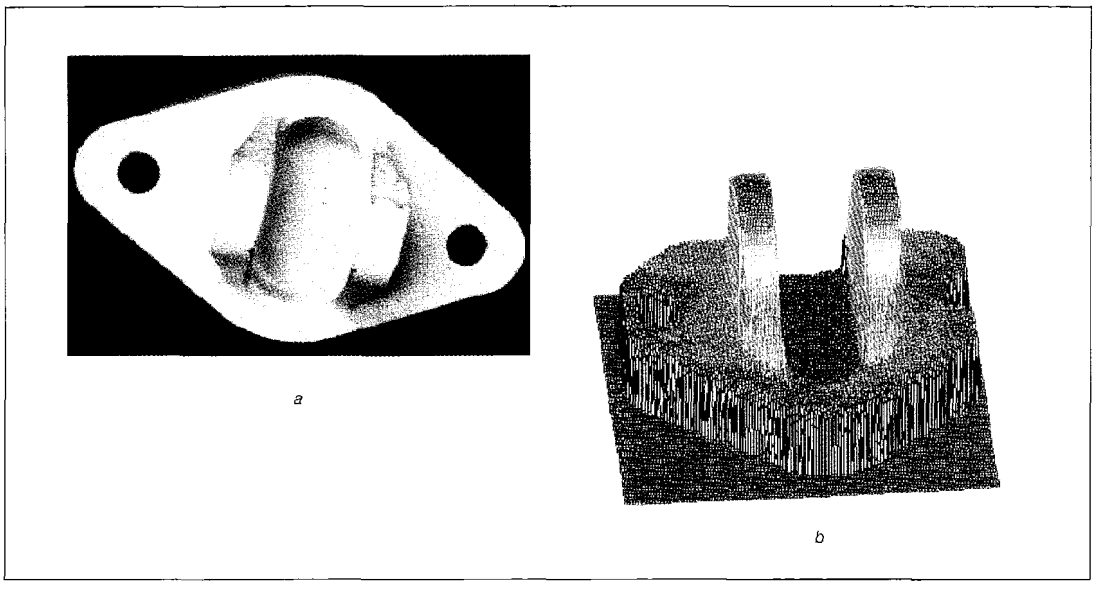

Fig. 10 Example of a small component; the door lock from an aeroplane is approximately 2cm across

room and temperature of the individual system components over extended periods of up to two days. The main sensitivity to variation in temperature is in the laser driver, but we have been able to correlate this exactly with the distance measurement. This makes temperature control of an engineered system a priority, although an alternative is to use a second known reference to compensate for the effects of long-term drift.

We are indebted to colleagues and friends within the University, at British Aerospace and at the National Museum of Scotland for giving us the opportunity to acquire 3D images of a number of test objects of varying size. In particular, we have made two extended visits to British Aerospace sites to take images of large and small components. Fig. 10 shows an example of a very small aerospace component, an aircraft door locking mechanism, which is approximately 2cm across. In this case we actually used a translation stage similar to that depicted in Fig. 1, rather than a pan-and-tilt head, as this gave even greater (x,y) accuracy.

The second example (Fig. 11) is of larger scale, about 40cm across. This comes from the National Museum and is a cast (around 1900 AD) of an original Roman frieze found near Jedburgh in the second century AD. The hole that is visible in the intensity image was actually used to hang it up! The depth data is shown in the form of a

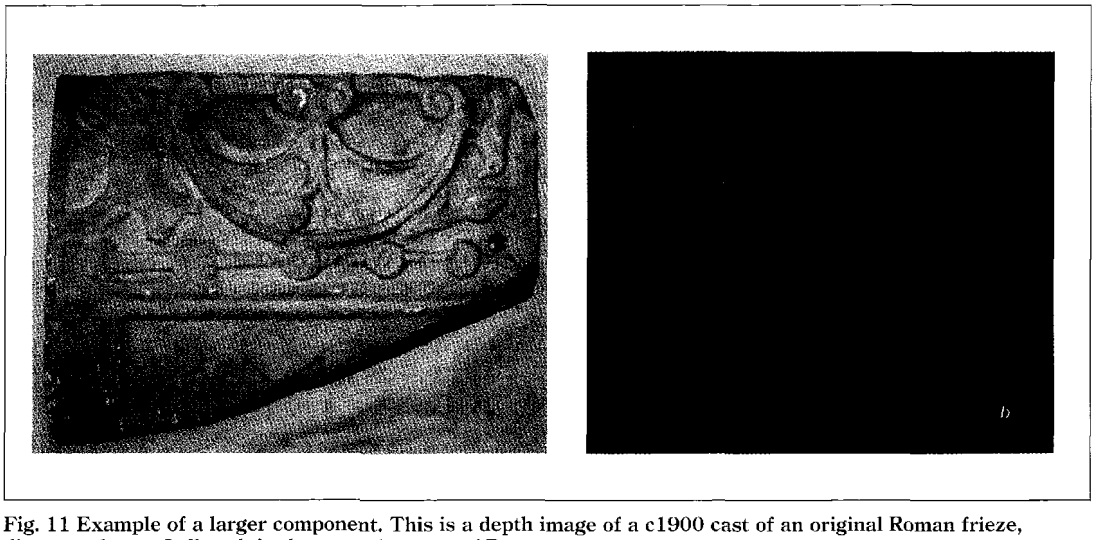

Fig. 11 Example of a larger component. This is a depth image of a c1900 cast of an original Roman frieze, discovered near Jedburgh in the second century AD

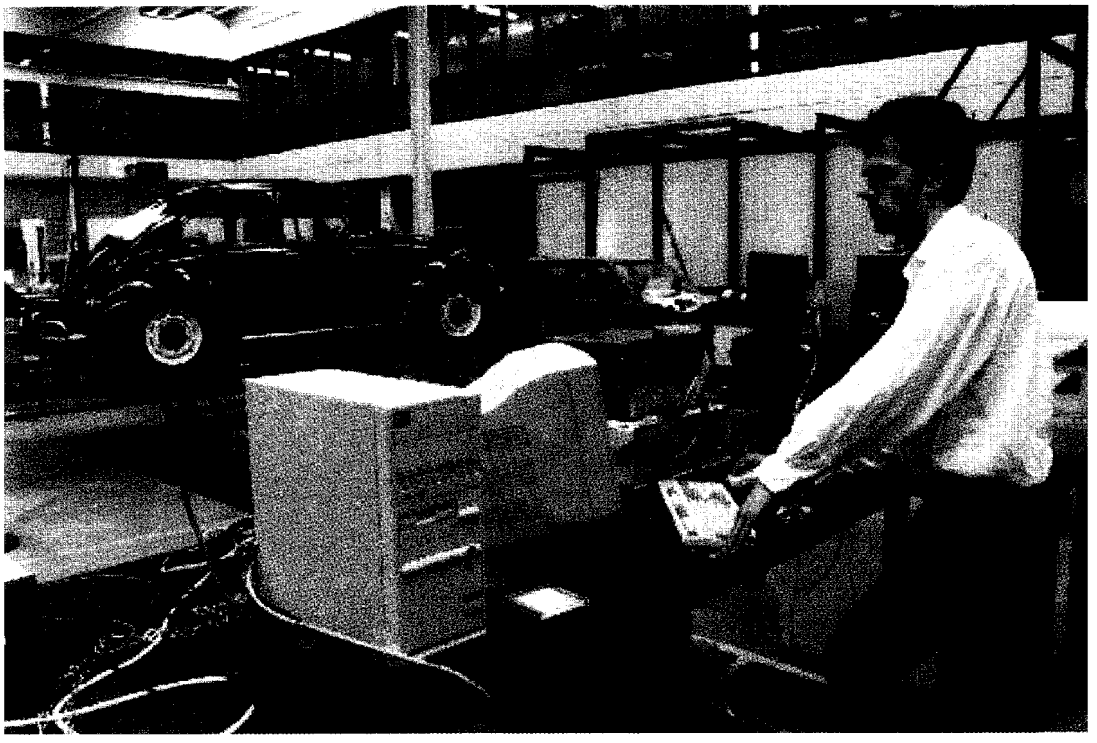

Fig. 12 Scanning a Volkswagen Beetle in the Mechanical Workshop

pseudocolour image, in which the 'hotter' colours are nearer to the sensor.

Our final example of a much larger object was taken in the Mechanical Engineering workshop at the university, as can be seen from Fig. 12. The sensor is visible mounted on the pan-and-tilt head just in front of the researcher. In this case the stand-off distance from the Volkswagen Beetle was about 4·5m as can be seen from the depth data displayed in pseudocolour form in Fig. 13. As the transverse spatial sampling was quite fine, the data file is very large indeed, so the displayed image is sub-sampled in this case for illustrative purposes.

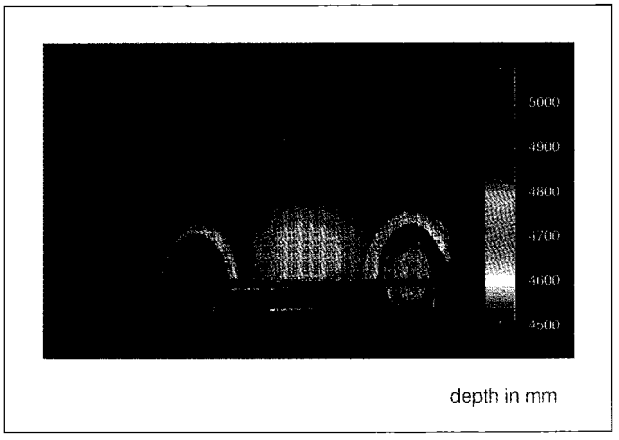

Fig. 13 The depth image of the Volkswagen, displayed as a pseudocolour image

Imaging with low photon counts: transparent and distant objects

Since time-correlated single-photon counting relies on the detection of the smallest quanta of optical energy, a system exploiting this technique is far more light sensitive than conventional analogue detection systems. This sensitivity has permitted measurements to be made of distant and uncooperative, for example transparent, targets, using both minimal collection optics and an eye-safe laser.

Fig. 14 shows an example of transparent object scanning; the object is a Lego piece with a transparent polygonal top, which is the surface of interest. The problem is not so much one of transparency (since ~4% of light is back-reflected from a glass surface and this is quite sufficient), but of specularity, since little of the reflected light may be directed back towards the detector. The two photon count histograms illustrate the problem. On the left is shown the return from the central surface, which was imaged at a near-specular direction. On the right is shown the return from one of the off-specular side faces,

which has a maximum count of only ten photons. As a consequence the longitudinal accuracy using existing methods is degraded in comparison with opaque objects and a scan of the object appears noisy.

On the other hand, Fig. 15 shows the photon return from a rectangular Perspex block in which near-normal incidence was maintained throughout. The three peaks on the left are the front and back surfaces of the block followed by a backplane; the peak on the right is the reference. In this case it is easy to measure the distance to, and the thickness of, the block. Thus, it is possible to measure several distances through several media simultaneously, not possible by previous techniques.

Expanding on this theme we have also conducted measurements of surfaces at much longer ranges of up to 50m.[8] We have reproduced the effects of even longer ranges by using neutral density filters to substantially reduce the photon return and still make measurements, effectively extending the range of the graph in Fig. 9 to obtain an estimate of depth resolution of 6mm with only ten photons collected. In comparison, a simulation of the photon data using Poisson statistics and modelling the shape of the histogram by piecewise exponential functions gave a theoretical estimate of 3mm, within a factor of two of the experimental measurement. For real distant measurements we set up the equipment in a long building, which houses an experimental wind tunnel, and obtained depth resolutions of 120μm at a distance of 50m with a 10s collection time, and of 1·5mm with a 1s collection time. Again, these were close to the results expected from the earlier tests.

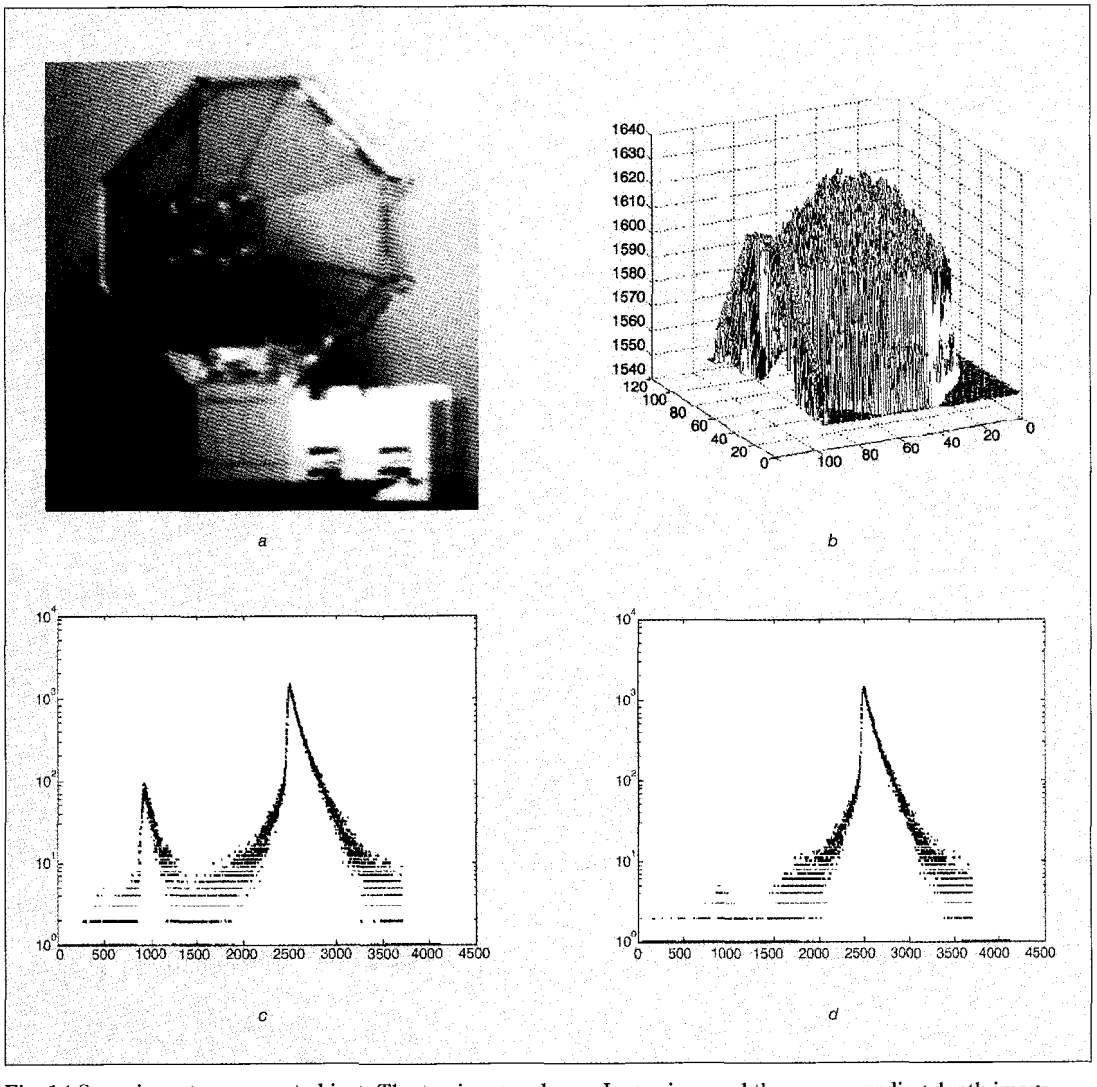

Fig. 14 Scanning a transparent object. The top images show a Lego piece and the corresponding depth image. The lower figures show examples of histogram returns from near and off specular directions

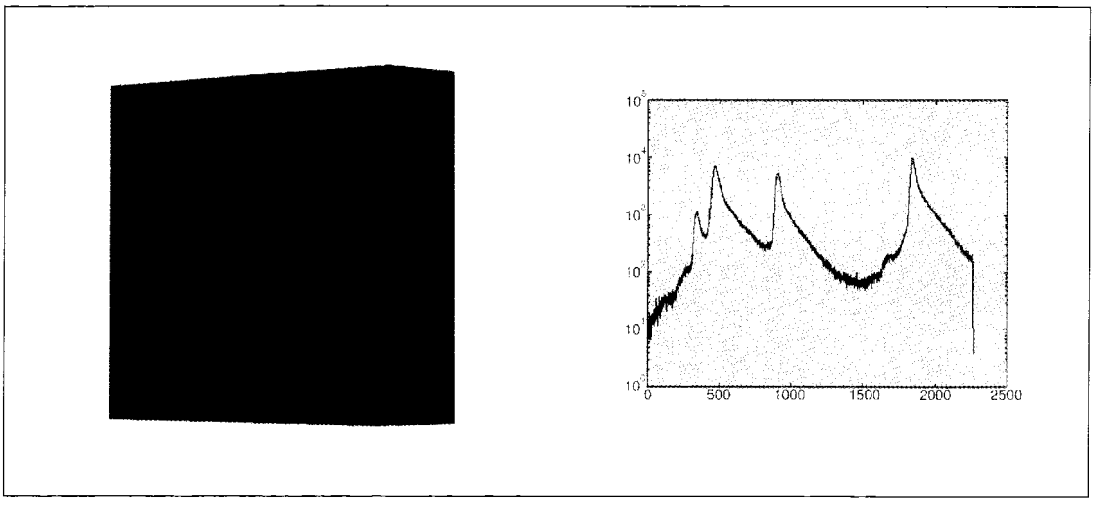

Fig. 15 Imaging a Perspex block. In this case the four peaks correspond, respectively, to the front and rear surfaces of the block, a cardboard back-plane and the reference channel

Conclusions

The objective of the work described here was to demonstrate the application of time correlated single photon counting to acquire accurate, occlusion-free, depth data from both cooperative and non-cooperative targets and ranging in size from a few centimetres to several metres. A summary of the existing sensor is as follows:

- Absolute measurement of range to an accuracy and repeatability of 15μm.
- Measurement in three dimensions using existing scanning to better than 0·1mm.
- Measurement of objects from less than 5cm to about 25m in size.
- Measurement directly from the object surface, not from targets or retro-reflectors.
- Robust and portable to allow imaging of large components *in situ*.
- Automatic surface scanning from many different surface materials.
- Eye-safe class 1 laser.

As an extension to the work we have also performed experiments on transparent objects and have assessed the performance of the system at distances beyond 50m. We believe that further improvements can be made to the optical design, the signal conditioning electronics and the method of processing the histogram data, especially at low photon counts, which would improve the performance of the technique still further.

Acknowledgments

We should like to acknowledge the contributions made to this work by J. Massa, M. Umasuthan, S. Pellegrini, G. Smith and P. Csakany. The bulk of the work was in collaboration with British Aerospace-Sowerby Research Centre and Edinburgh Instruments Ltd., supported financially by the UK Engineering and Physical Sciences Research Council. The passively Q-switched lasers are provided by E.L Portnoi and co-workers, A.F. Ioffe Institute, St. Petersburg. The SPAD diodes are provided by S. Cova and co-workers at the Politecno di Milano. We are also grateful to A. Sheridan at the National Museum of Scotland for allowing us to scan some rare artefacts.

References

1 CIPOLLA, R., and ROBERTSON, D. P.: '3D models of architectural scenes from uncalibrated images and vanishing points', Proceedings of the 10th IAPR International Conference on Image Analysis and Processing, 1999, pp. 824-829
2 CLARKE, T.: 'Photogrammetry finally makes it off the shelf', pp. 23-28, 'Simple scanners reveal shape, size and texture', *Optics and Lasers Europe*, May 1998, pp. 29-32
3 PALOJARVI, P., MAATTA, K., and KOSTAMOVAARA, J.: 'Integrated time of flight laser radar', *IEEE Transactions on Instrumentation and Measurement*, 1997, **46**, (4), pp. 996-999
4 KACYRA, B., DIMSDALE, J., and BRUNKHART, M.: 'Integrated system for quickly and accurately imaging and modelling three dimensional objects', Cyra Technologies Inc., US Patent 5,988,862, November 1999
5 GAGNON, E.: 'Laser range imaging using the self-mixing effect in a laser diode', *IEEE Transactions on Instrumentation and Measurement*, 1997, **48**, (3), pp. 693-699
6 MASSA, J. S., BULLER, G. S., WALKER, A. C., COVA, A. S., UMASUTHAN, M., and WALLACE, A. M.: 'A time-of-flight optical ranging system using time-correlated single photon counting', *Applied Optics*, 1998, **37**, (31), pp. 7298-7304
7 UMASUTHAN, M., WALLACE, A. M., MASSA, J., BULLER, G. S., and WALKER, A. C.: 'Processing time-correlated single photon data to acquire range images', *IEE Proceedings: Vision, Image and Signal Processing*, 1998, **145**, (4), pp. 237-243
8 PELLEGRINI, S., BULLER, G. S., SMITH, J. M., Wallace, A. M., COVA, S.: 'Laser based distance measurement using picosecond resolution time correlated single photon counting', *Measurement Science and Technology*, 2000, **11**, pp. 712-716

The authors are with the Departments of Computing and Electrical Engineering, and Physics, Heriot-Watt University, Edinburgh EH14 4AS, UK.