33rd Annual International Conference of the IEEE EMBS
Boston, Massachusetts USA, August 30 - September 3, 2011

Next Generation Optical Surface Sensing for Real-time Measurement in Radiotherapy

James M. Parkhurst, Gareth J. Price, Phil J. Sharrock, Tom E. Marchant and Christopher J. Moore

***Abstract*—With the introduction of intensive new treatments such as hypo-fractionation and proton beam therapy, localization of the tumor target volume and tracking of points across the skin entrance surface have become critically important. Optical metrology has been used to monitor the patient's bulk position and motion throughout treatment. However systems have not been capable of high temporal and spatial resolution whole-surface topology measurement. We describe the implementation of such a system based on Fourier profilometry. Its algorithm is split into four separate processing stages, including spatial phase determination: descriptions of each stage are given along with the modifications made to increase performance. The optimized system is capable of processing 23 frames per second (fps), with each frame providing 512x512 measured points. The data density, accuracy and performance of the system are an order of magnitude improvement on commercially available clinical systems. We show that this performance permits genuinely real-time measurement of a patient, live during both setup and radiation treatment delivery. It is also fast enough to provide smooth dynamic visualizations of motion at all points on the wrap-around body surface for radiotherapy staff and intuitive, direct feed-back to patients.**

I. Introduction

In recent years, new intensive radiation therapies, such as hypo-fractionation and proton beam therapy, have become available with the potential to greatly improve patient outcomes. These exemplars are associated with an increased need to target dose precisely, with motion and shape-changes being confounding factors. With scanning beam proton therapy in particular, changes in surface topology can have an appreciable effect on the dose received by the tumor and the surrounding healthy tissue. Optical non-contact measurement of the patient's body surface has the potential to measure the changing skin-topology. However, these computationally demanding techniques have only delivered low frame rates rather than the genuinely real-time performance required within the clinic.

One approach to measuring moving skin topology is to use point by point scanning. Recently a laser triangulation spot system intended for skin range measurement during proton beam therapy treatment has been simulated [1,2]. The system measures the change in height at a single point to model the motion of the beam entrance surface in an attempt to compensate for its effects on treatment delivery. Stereo-photogrammetric methods have been used to extrapolate a larger area of the patient's body from a sparser, underlying measured point cloud that is characteristically irregular. Full surface performance of up to 2Hz has been reported [3], which can only be improved upon by selecting a smaller region of interest. The accuracy of the underlying measured points in particular and extrapolated surfaces in general have not been reported, with studies instead focusing on bulk displacement.

A dynamic optical sensor system, developed at The Christie Hospital, has been utilized for offline monitoring and analysis of patients undergoing radiotherapy [4-10]. This is supporting the development and clinical deployment of a device implementing real-time feedback of patient position and pose in support of improved immobilization during irradiation [11]. Being based on Fourier profilometry, it makes use of cosinusoidal structured light projection, where the spatial phase modulation of the projected pattern is used to determine the coordinates of all illuminated points in each video frame of captured data. Accordingly, it can also be used for 'real-time' control, provided the combined projection, video-capture, computational processing and output become fast enough: this is commonly defined in the context of visualization applications where rates of 25 Hz are considered to be genuinely real-time. Henceforth we use this definition of real-time for surface topology measurement.

We present a genuinely real-time optical sensor system capable of showing the patient's skin-topology in the treatment room 23 times per second, with each frame providing a regular array of 512x512 3D measurement points. In its current operational configuration it is positioned at a distance of approximately 3m obliquely above the patient, as dictated by the treatment machinery in the radiotherapy bunker. Nevertheless, the accuracy of the measured points, within a 40cm^3 volume calibrated at patient level, is to within 2mm of an isosurface extracted from a simultaneously captured cone beam CT with a significant

Manuscript received June 17, 2011. This work is supported by UK grants from the Engineering and Physical Sciences Research Council (EP/D078415/1) and the National Institute for Health Research (NEAT K021).

J. M. Parkhurst is with North Western Medical Physics at The Christie NHS Foundation Trust, Manchester, M20 4BX, UK (phone: 0161 446 3258; email: james.parkhurst@physics.cr.man.ac.uk).
G. J. Price is with North Western Medical Physics at The Christie NHS Foundation Trust (email: gareth.price@physics.cr.man.ac.uk).
T. E. Marchant is with North Western Medical Physics at The Christie NHS Foundation Trust (email: tom.marchant@physics.cr.man.ac.uk).
P. J. Sharrock is with North Western Medical Physics at The Christie NHS Foundation Trust (email: phil.sharrock@physics.cr.man.ac.uk).
C. J. Moore is with North Western Medical Physics at The Christie NHS Foundation Trust (email: chris.moore@physics.cr.man.ac.uk).

proportion of points achieving 1mm accuracy. In this paper we focus on the algorithmic development that was required to achieve this performance.

II. Method and Materials

The optical sensor system we describe is a fringe projection system. Structured light in the form of a cosine fringe pattern is projected on to the object to be measured – in this case a human patient undergoing radiotherapy. The 3D surface is generated by analysis of the modulation of the fringe pattern across the object with respect to a flat reference plane using the Fourier profilometry technique [12].

Severe constraints are put on the geometry of the optical sensor system by the typical configuration of equipment already in the radiotherapy treatment room such as the linear accelerator gantry itself. The practical outcome of this is that the cameras are setup at a distance from and at an oblique angle to the patient. This causes shadowing in regions that aren't visible to the camera and limits the area of the patient's body that can be measured using a single sensor system. To overcome this problem, the system described here uses three color separated sensors as described in [13]. Although this method of separation allows the use of multiple sensors, it comes at the expense of a significant reduction in signal-to-noise ratio. Fig. 1 shows the geometry of the system relative to the linear accelerator gantry and treatment couch. The central green sensor is setup such that it is always possible to measure the patient; the left and right sensors are setup to provide the wrap-around capability; however they can become partially occluded due to the rotation of the gantry into the line of sight.

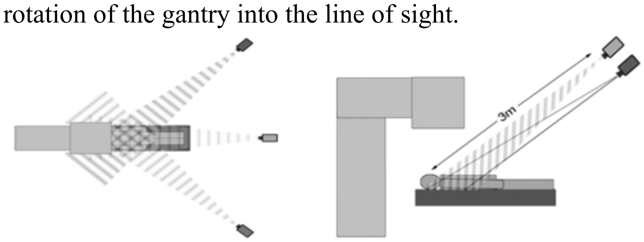

(a) (b)

Fig. 1. A diagram showing the layout of the optical sensors heads relative to the treatment machine and couch (left hand grey blocks) from above (a) and from the side (b). The three camera/projector pairs are separated using red, green and blue narrow band filters and each capture a portion of the patient body surface, the partial surfaces combine to create a full 'wrap-around' surface.

The original source code for the optical sensor system was implemented in a mixture of C/C++ and ITT's Interactive Data Language (IDL)[1]. This prototype system achieved high surface data densities of 512x512 measured points per frame but, in our view, was primarily suitable for off-line processing of patient data, since it achieved a modest temporal resolution of 1-2 fps. It was therefore necessary to radically redesign the system to achieve true real-time performance; the source code was rewritten using the C programming language, high performance image processing libraries were used and micro-optimizations were applied to the computationally intensive parts of the algorithm. Our implementation of the Fourier profilometry algorithm can be split into the following four processing steps: calculation of the reference spot coordinates; calculation of the wrapped phase map from the fringe image; calculation of the unwrapped phase map; and conversion from the relative phase values to the calibrated real-world coordinates. A description of each of these is given below.

A. Calculation of Reference Spot Coordinates

Reference spots are used to identify the location of a specific fringe in each fringe image; this enables the relative phase with respect to a flat reference plane to be calculated for use in the height calibration. Though there are various pitfalls to try and avoid in the accurate detection of spots in images, the general algorithm employed here is very simple, easy to implement and computationally efficient. The spots are isolated from the background by applying the Laplacian of Gaussian operation to the fringe image. Since the spot size is found during calibration of the system, multi-scale operations are not required. The spots are then identified by searching along the track traversed by each spot across the CCD acquired during calibration of the system. As the spot tracks are well-defined and by design do not overlap at any point, each can be easily identified from frame to frame. The computation time required for the calculation of the coordinates of the seven reference spots we use (for reasons of redundancy) is ~5ms.

B. Calculation of Wrapped Phase Map

The full algorithm used to calculate the wrapped phase map is described in [10]. Briefly, the processing can be summarized as follows. The image is first filtered in order to removing the background leaving the fringes as the major component. The Fourier transform of the image is then calculated and the image filtered in the frequency domain using a band-pass filter to isolate the peak corresponding to the frequency of the fringes. Upon being transformed back to the spatial domain, the wrapped phase map is calculated as the argument of the complex image.

Fast algorithms exist for the calculation of Fourier transforms and for performing various image processing functions. Many high performance commercial and open source software libraries are generally available for performing these tasks. Our software has been written to target the X86 processor architecture. Since we make use of Intel processors in our system, we use the Intel Integrated Performance Primitives (IPP) library[2] optimized for use on Intel processors. By utilizing the IPP library and applying micro-optimizations to the image processing source code, the computation time required for the calculation of the

[1] ITT, Boulder, USA

[2] Intel Corporation, Santa Clara, USA

wrapped phase map from the input fringe image was reduced to ~15ms.

C. Calculation of Unwrapped Phase Map

The wrapped phase map is modulo 2π; the values are wrapped between π and $-\pi$. In order to find the true phase values across the field of view of the instrument a phase unwrapping algorithm needs to be employed. At the time of writing, there are no generally available software libraries commercial or otherwise that implement a selection of general purpose phase unwrapping algorithms; therefore, it is either necessary to use an available research implementation or to implement an algorithm directly from its published description. An investigation was conducted in order to identify which phase unwrapping algorithms were best suited for use in a real-time Fourier profilometry system [14]. Algorithms were assessed based on their speed of execution and their ability to 'correctly' unwrap the given phase maps generated from data taken from the free-form shape of real human subjects using the optical sensor described above. A great number of phase unwrapping algorithms exist in the literature; however, only those with a readily available implementation were used in the investigation [15,16].

As a result of this investigation, it was found that despite the use of modern high performance computer hardware, none of the algorithms evaluated were capable of the performance required for use in a real-time system whilst maintaining accurate results. Execution times for the algorithms ranged from between 13ms and 5s for the datasets used; however as expected there was a trade off between speed of execution and robustness of result. It was therefore necessary to make modifications to an existing algorithm - the non-continuous quality-guided path algorithm described in [17] - in order meet the specified requirements of the system. These modifications were made for the sole purpose of improving the execution time of the algorithm; there were no observable differences in the resulting unwrapped phase maps with those from the unmodified algorithm. The resulting modified algorithm introduced a parameter-free preprocessing step to partially unwrap the phase map before applying the unmodified algorithm to complete the unwrapping process. Using the modified algorithm, a full 512x512 pixel phase map can be unwrapped in ~30ms. Using the mask generated by applying a threshold to the magnitude of the complex fringe image typically leaves around 100,000 points where an accurate measurement can be made. In this case, the algorithm is able to unwrap the masked phase map in ~9ms.

D. Conversion to Real-World Coordinates

In order to convert the relative phase map into a full 3D surface, three simple equations, whose parameters are found during calibration of the system, are applied to each of the phase values. The height (y coordinate) at each point is calculated by application of a non-linear equation. The x and z coordinates are then calculated by applying a quadratic perspective correction using the calculated heights. The computation time required for this step, ~1ms, is negligible compared to that of finding the spots and calculating the unwrapped phase map.

III. Results and Discussion

Following the modification and optimization of the Fourier profilometry software, the execution time for each stage of the surface calculation was greatly improved. Table I shows the executions times of both the original software and the optimized software. At each stage of the calculation the improvement can be seen; in some cases the improvement is over an order of magnitude in execution time. For measured data arrays of 512x512 3D points, from raw video frame acquisition through computation of the real-world Cartesian coordinates of points to their visual display, the original software was capable of 1 - 1.5 fps for a single data channel; the optimized software is capable of ~23 fps. If all three color-separated data channels are used, the optimized software is currently capable of performing acquisition to display routinely at a rate of ~16 fps.

TABLE I
Execution Times for Original and Optimized Software

Processing Step	Computation Time	
	Original	Optimized
Spot Coordinates	190ms	5ms
Wrapped Phase	125ms	15ms
Unwrapped Phase	140ms	9ms
Real-world Coordinates	< 1ms	< 1ms

The main driver of visible anatomical motion in radiotherapy patients is considered to be respiration. In practice, respiratory motion is highly variable but is commonly described as having a period of approximately 4 seconds for healthy individuals. It can be much faster, e.g. 1 second, when subjects are anxious and particularly ill, as with cancer patients undergoing concurrent radiotherapy and chemotherapy. However, discomfort during treatment delivery, which typically lasts a quarter of an hour or more, can be more significant. Our studies have shown that it can result in large, complex pose changes [8]. Sudden, intra-fraction transients due to discomfort or relaxation, if detected by observant clinical staff, can warrant corrective action for photon treatments, with more severe consequences in proton therapy. Clearly there is an increasing need for high speed surface measurement.

The real-time capability of our system has been utilized for live measurement and visual feedback of the patient's body surface position and motion during setup and treatment. The purpose of providing visual feedback of the patient's motion is to be able to provide additional support for both patient and radiographer in order to help the patient keep still within certain bounds during radical treatment. Fig. 2 shows a surface produced by the optical sensor (a) and visualization from the feedback device (b). The motion is visualized as the difference between the dynamic measured body surface and the planned treatment position; bounds are given in order to indicate the range of motion deemed to be acceptable. Taking the difference between the live patient surface and the planned treatment position results in a 'lamina' that flexes with the patient's breathing motion and forms an effective and intuitive visualization of the range of motion experienced across the whole region of interest on the patient's body.

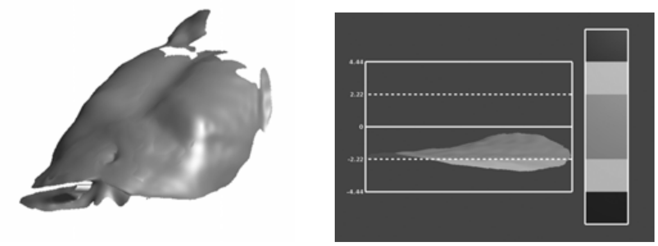

(a) (b)

Fig. 2. A full wrap-around patient surface produced by the optical sensor (a) and a visualization showing the patient's motion in a region of interest relative to the planned position for the central surface (b); threshold bars and colors are used to indicate the range of acceptable motion. The spacing of the threshold bars, calculated from the motion of the patient, is 2.22 mm.

In order to provide an effective visualization that displays the patient's motion smoothly and intuitively, video frame rate performance is a necessity. Lower frame-rates could result in a visualization that appears discontinuous; systems with frame rates around 1 fps will be unable to measure transient body shifts or portions of a rapid breathing cycle, thereby severely limiting the usefulness of visualization. Speed is especially important when providing visual support for patients since the smoother the visualization appears; the easier it will be for them to interpret their motion. Efficient algorithms for analysis of the surface data, applied as a post-processing step to the surface creation algorithm, mean that the motion and analysis can be simultaneously displayed at the frame rates required for this application. The computer system on which the optical sensor software is run uses dual Intel Xeon processors containing a total of 12 processor cores and 24 hardware threads (when hyper-threading is used). A single sensor channel is capable of running at full video frame rate using a single processor core, therefore with refinements to better utilize resources it is anticipated that 3D reconstructive frame rates very close to 25Hz will be achieved for the three sensor system.

The accuracy of the measurements produced by the optical sensor system is to within 2mm for 90% of all measured points. This is comparable with the accuracy achieved by the cone beam CT scanner with which the measurements were validated. This accuracy certainly matches the 2mm tolerances set for typical treatment machine isocentre reproducibility. With a significant proportion of points approaching 1mm accuracy, skin surface tracking with the spatial and temporal resolution required for proton therapy appears to be feasible.

IV. Conclusion

We have described significant improvements to the algorithms underlying an optical body surface sensor using the technique of Fourier profilometry, resulting in genuine real-time performance under clinical conditions. Measurement of 512x512 3D points on the patient's body surface proceeds at a rate up to 23 fps with a spatial accuracy consistently better than 2mm, which matches the capabilities of treatment linear accelerators and in room X-ray imaging.

References

[1] N. Saito, et. al., "Speed and accuracy of a beam tracking system for treatment of moving targets with scanned ion beams," Phys. Med. Biol., vol. 54, no. 16, pp. 4849-4862, July 2009.
[2] C. Bert, et. al., "Dosimetric precision of an ion beam tracking system," Rad. Onc. vol. 5, no. 61, June 2010.
[3] J. L. Peng, et. al., "Characterization of a real-time surface image-guided stereotactic positioning system", Med. Phys., vol. 37, no. 10, pp. 5421-5433, Oct. 2010.
[4] F. Lilley, M. J. Lalor, and D. R. Burton, "Robust fringe analysis system for human body shape measurement," Opt. Eng., vol. 39, pp. 187, no. 1, Jan. 2000.
[5] C. J. Moore, et. al., "Opto-electronic sensing of body surface topology changes during radiotherapy for rectal cancer," Int. Jnl. Rad. Onc. Biol. Phys., vol. 56, no. 1, pp. 248-258, May. 2003.
[6] G. J. Price, et. al., "An analysis of breast motion using high-frequency, dense surface points captured by an optical sensor during radiotherapy treatment delivery," Phys. Med. Biol., vol. 54, no.21, pp. 6515-6533, Nov. 2009.
[7] G. J. Price, et. al., "Design of, and some clinical experience with, a novel optical surface measurement system in radiotherapy," in Proc. SPIE. Medical Imaging, San Diego, 2010, 76224B.
[8] C. J. Moore, G. Whitfield, P. Price, and J. M. Parkhurst, "Patient pose changes during radiotherapy measured using 3D colour Fourier profilometry," Intl. Jnl. Comp. Assist. Radiolog. Surg., vol. 5, no. 1, pp. 70-71, June 2010.
[9] G. J. Price, et. al., "4D cone beam CT phase sorting using high frequency optical surface measurement during image guided radiotherapy," in Proc. SPIE. Medical Imaging, Florida, 2011, 79615S.
[10] G. J. Price, et. al., "A real-time robust Fourier profilometry system for use in image guided radiotherapy," unpublished.
[11] J. M. Parkhurst, et. al., "Visualisation of Dynamic Surface Data for a Patient Display to Reduce Movement during Radiotherapy", in Proc. Biomedical Engineering, vol. 2, Innsbruck, 2010, 680-074.
[12] M. Takeda, and K. Mutoh, "Fourier-Transform Profilometry for the Automatic-Measurement of 3-D Object Shapes," App. Opt., vol. 22, no. 24, pp. 3977-3982, Dec. 1983.
[13] O. A. Skydan, M. J. Lalor, and D. R. Burton, "Technique for phase measurement and surface reconstruction by use of colored structured light," App. Opt., vol. 41, no. 29, pp. 6104-6117, 2002
[14] J. M. Parkhurst, G. J. Price, P. J. Sharrock, and C. J. Moore, " Phase unwrapping algorithms for use in a true real-time optical body sensor system for use during radiotherapy," App. Opt., submitted for publication.
[15] D. C. Ghiglia, and M. D. Pritt, "Two-Dimensional Phase Unwrapping: Theory, Algorithms, and Software," New York, Wiley, 1998.
[16] M. A. Herráez. (2008, Oct 6). Two-Dimensional Phase Unwrapping. [Online]. Available: http://www.ljmu.ac.uk/GERI/90207.htm
[17] M. A. Herráez, D. R. Burton, M. J. Lalor, and M. A. Gdeisat, "Fast two-dimensional phase-unwrapping algorithm based on sorting by reliability following a noncontinuous path," App. Opt., vol. 41, no. 35, pp. 7437-7444, Dec. 2002.