Current status of the CBCT project at Varian Medical Systems

Evangelos Matsinos

Varian Medical Systems Imaging Laboratory GmbH, Täfernstrasse 7, CH-5405 Baden-Dättwil, Switzerland (evangelos.matsinos@varian.com)

Abstract

Tracking and targeting the tumours are simultaneous processes in the image-guided radiotherapy (IGRT); this is expected to boost the efficiency, the reliability, and the safety in the treatment. Varian Medical Systems (VMS) has already produced and installed the first IGRT machine; the device comprises the VMS Clinac equipped with the On-Board Imager (OBI) component. Cone-beam CT (CBCT) imaging, one of the options of the OBI machine, aims at high-quality volumetric reconstruction. A number of calibrations are needed in order to operate our CT-imaging machines properly; they ensure that the machine components are properly aligned, the mechanical distortions are small, and yield important output that is used in the reconstruction of the actual scan data. The geometrical calibration is achieved by using a needle phantom. In order to increase the dynamic range of our imager (hence, to obtain reliable information simultaneously in the high- and the low-attenuation areas of the irradiated object), VMS has developed a dual-gain mode. Next on our agenda is the suppression of (ring, streak, and beam-hardening) artefacts in our reconstructed images and the further development of our detectors in order to remove patterns relating to lag and ghosting effects.

Introduction

Accurate knowledge on the position of the tumour during the treatment has long been one of the harassing issues in radiotherapy. To overcome the absence of a monitoring process, practical enough to apply in parallel to the actual treatment or even on a daily basis, the oncologists 'traditionally' augmented the target area in order to maximise the probability of destroying the entire tumour by the end of the treatment. The inevitable consequence of the enhancement of the region, receiving the largest radiation impact, is the destruction of healthy tissue.

In the war against cancer, Varian® Medical Systems (VMS®) has recently deployed a number of new products aiming at increased precision in radiotherapy. Last year, our armoury was enriched with the VMS Clinac® coupled with the On-Board Imager™ (OBI) system (Fig. 1). This device became the first clinically-applicable solution for real-time, image-guided radiotherapy (IGRT); it enables tracking and targeting the tumours 'simultaneously'. Evidently, the OBI component facilitates the delivery of higher dose in the target and ensures better protection of the surrounding

Medical Imaging 2005: Physics of Medical Imaging, edited by Michael J. Flynn,
Proceedings of SPIE Vol. 5745 (SPIE, Bellingham, WA, 2005)
1605-7422/05/$15 · doi: 10.1117/12.593418

healthy tissue[1].

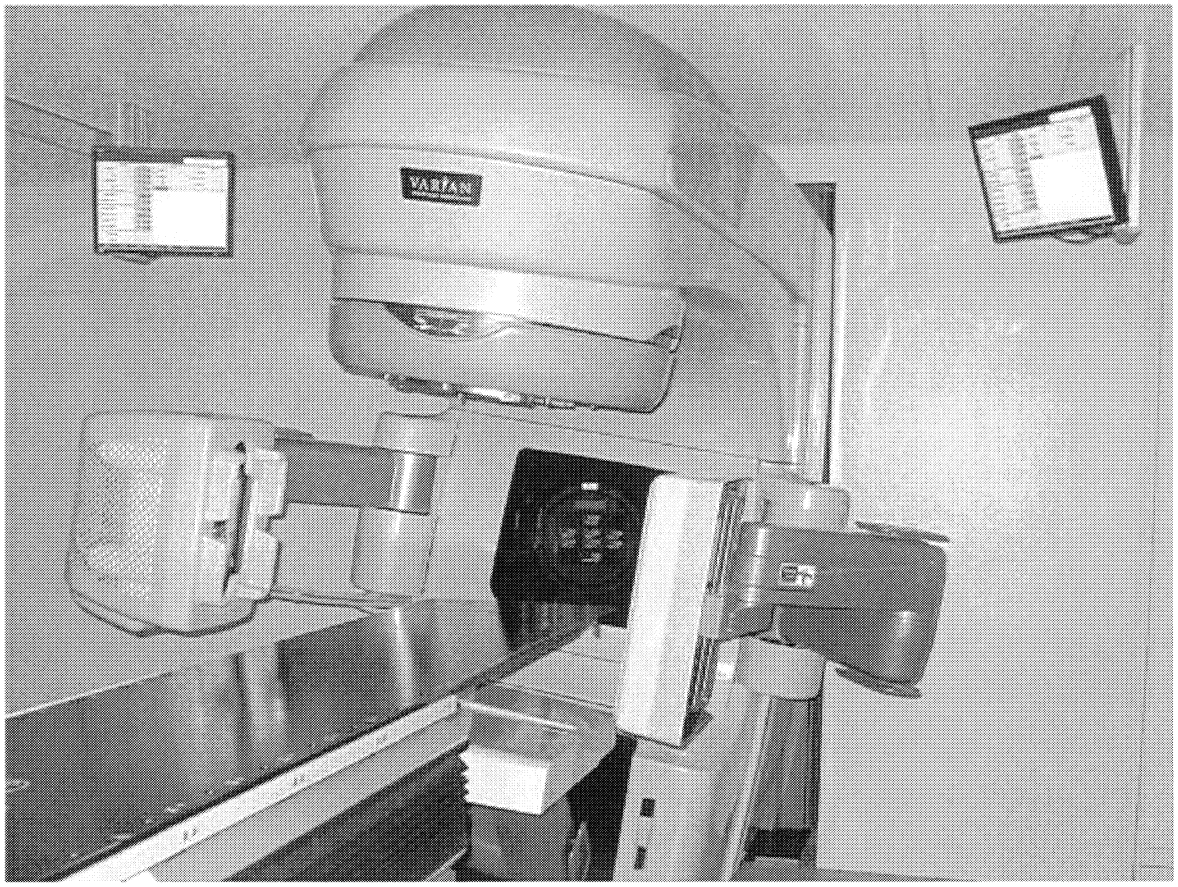

Figure 1. The Clinac featuring the On-Board Imager module; Hirslanden Klinik, Aarau, Switzerland

The two main components of the OBI are the low-energy (kV) X-ray tube (presently, the VMS G242 model) and the PaxScan® 4030CB flat-panel detector (FPD); the FPD is attached to the body via a system of robotic arms enabling 3D movement. Our FPDs are capable of doing fast fluoroscopy, to be used for position and verification (real-time moving images), as well as superior radiography to be used for diagnosis (high-resolution images); relating to the latter is our cone-beam CT (CBCT) project. The aim of the present paper is to provide information on the current status of the CBCT project; this progress report relates both to our Acuity™ Imaging System (which is a dedicated simulation and CT-imaging device, see Fig. 2) and to the OBI.

---

[1] Up to now, the OBI has been used to determine the shift in the tumour location prior to the treatment session (interfraction shift). The 'online' assessment (Dynamic Targeting™) of the tumour position (intrafraction shift), contemporaneous with the dose delivery in a scheme that encompasses Real-Time Position Management™ (RPM™) respiratory gating, is expected to boost radiotherapy in the years to come.

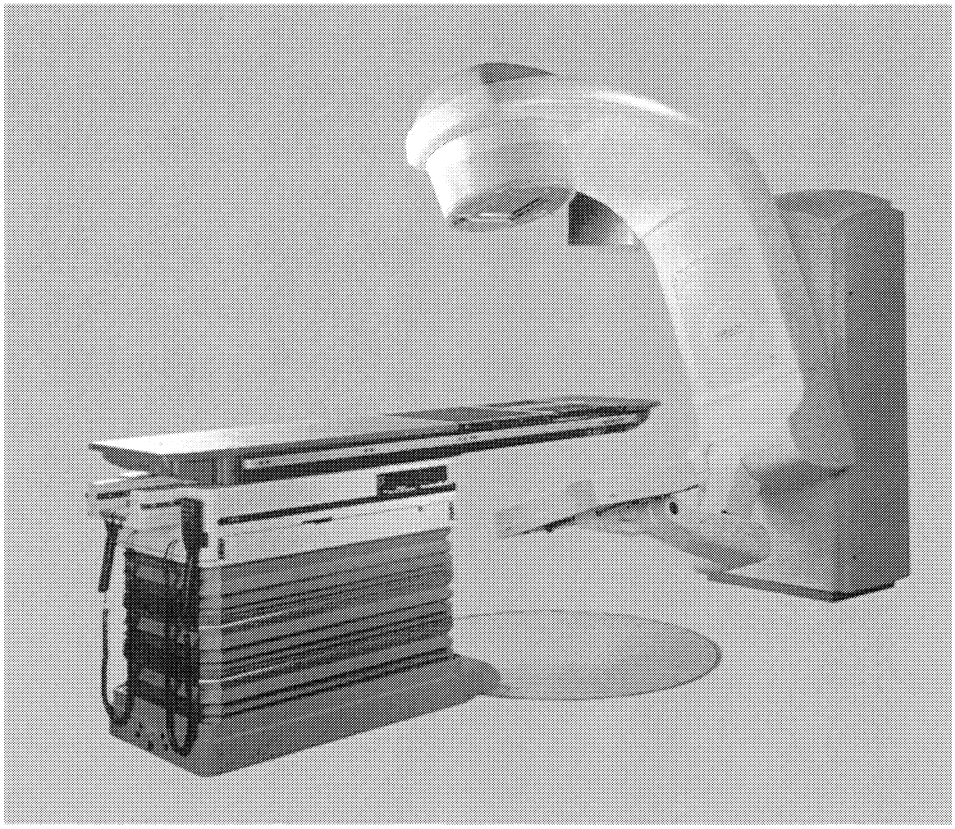

Figure 2. The Acuity Imaging System

Our 4030CB amorphous-silicon FPD is a real-time digital X-ray imaging device consisting of an array of 2,048X1,536 square pixels (side length: 0.194mm) and spanning an approximate area of 40X30cm^2. In order to expedite the data transfer and their processing, the so-called half-resolution (2X2-binning) mode is used in our applications; thus, the imager is assumed to consist of 1,024X768 (logical) square pixels with a side length of 0.388mm. Due to the high sensitivity of the scintillating material (CsI) and to sophisticated noise-reduction techniques, the low-dose imaging performance of our FPDs is remarkable.

Calibrations

To obtain reliable high-quality volumetric information, the CT device must be properly calibrated; the calibrations may be classified as general and mode-specific.

a) The general calibrations include:
   - the geometrical calibration,
   - the I_0 calibration,
   - the norm-phantom calibration,
   - the calibration for the beam-hardening correction, etc.

b) The mode-specific calibrations relate to the different operating modes of the imager.

The geometrical calibration

The purpose of the geometrical calibration of the CT apparatus is threefold.

- To ensure that the machine components are properly aligned.
- To provide the imager lateral and longitudinal offsets (including their dependence on the gantry angle) that are important in the reconstruction[2].
- To monitor the mechanical stability of the machine.

Utilised in the geometrical calibration is a needle phantom comprising five 6cm-long 3mm-in-diameter metallic needles embedded in a urethane ($C_3H_7NO_2$) compound. It is mounted directly onto the table and is properly aligned with the help of a laser system. The needles are placed along the longitudinal direction (i.e., along the table). The couch angles (azimuth and tilt) are set to 0^0. The input data comprise one complete scan (360^0 rotation of the gantry). One typical image, obtained around 0^0, is shown in Fig. 3.

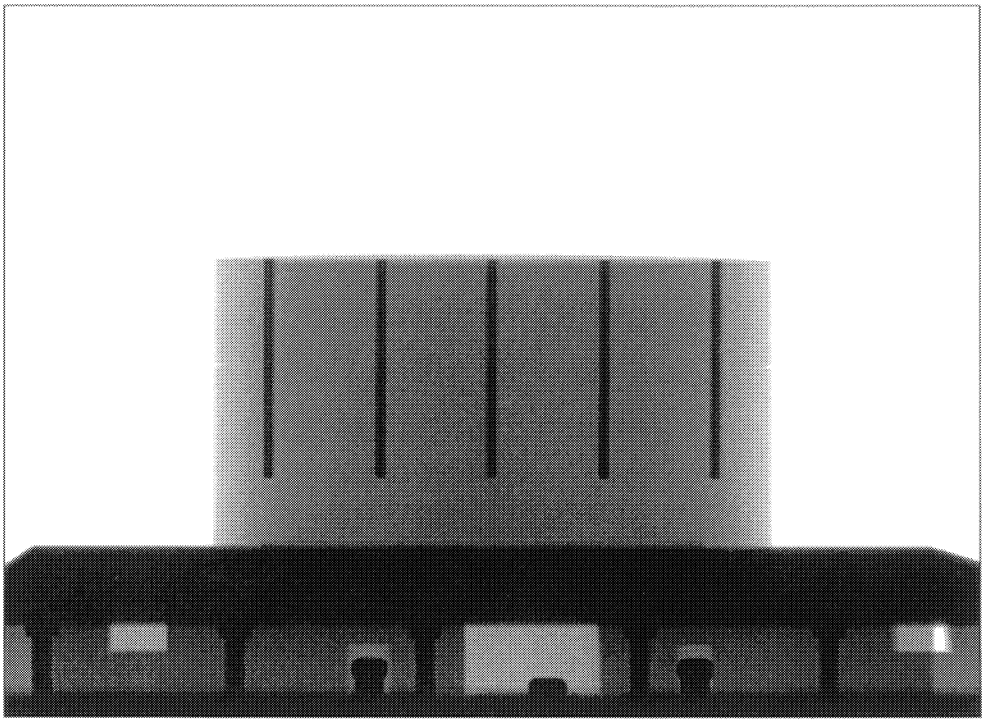

Figure 3. A typical image of the needle phantom

[2] In case that these offsets are not supplied at reconstruction time, the resulting image may be fuzzy (i.e., lacking sharpness).

The average signal along the imager (i.e., along its 40cm-long side) is firstly constructed from 16 rows of pixels neighbouring the midline (8 on either side). To get rid of the background, produced by the urethane case, a fixed-width window runs over the data, summing up the contents and subtracting a linear background (as defined by the two outmost pixels). The window width has been chosen slightly larger than the diameter of the needles projected onto the imager. Applying this approach to the data yields the transformed signal, shown in Fig. 4, which is easy to analyse; the negative peaks correspond to the axis positions of the needles. The air-urethane borders have been eliminated via a simple software cut.

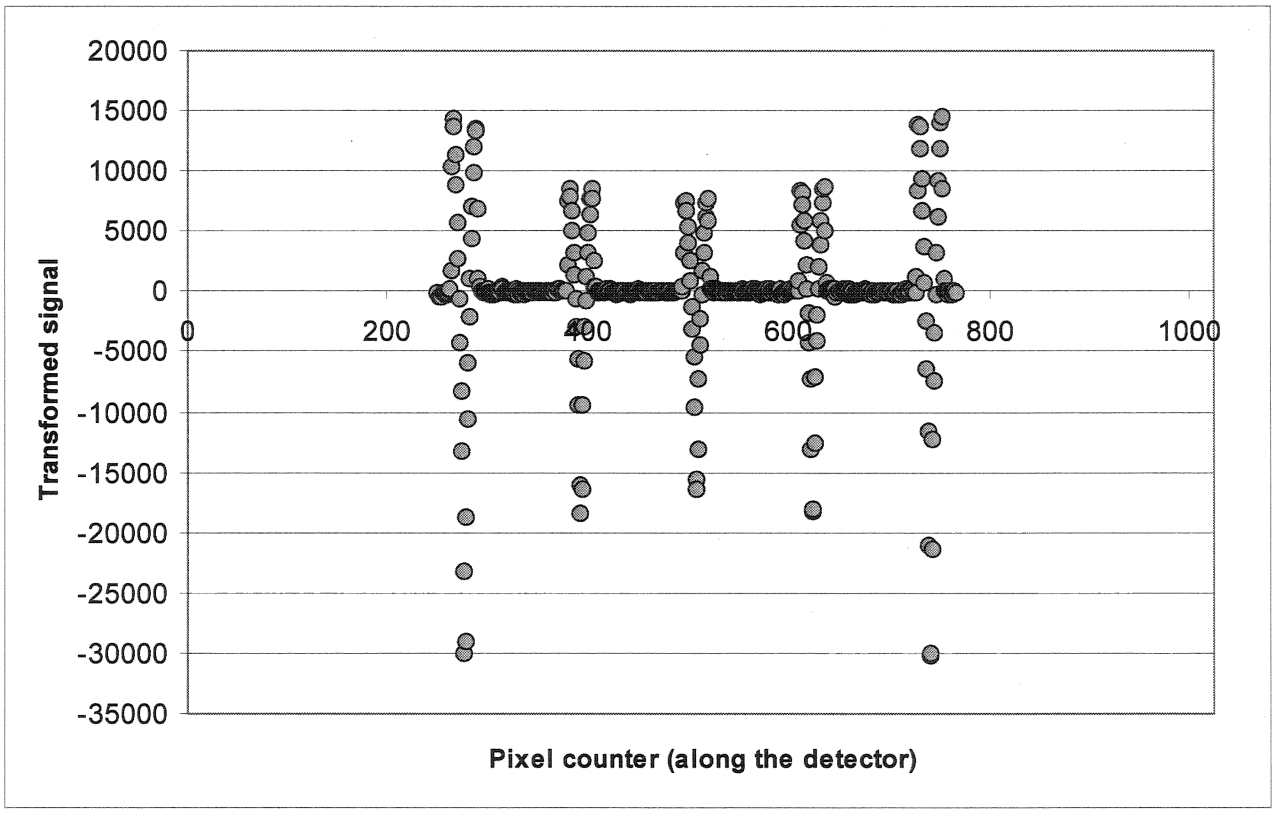

Figure 4. A typical example of the transformed signal along the imager

After the signal modes have been properly assigned to the needles, each needle is followed towards the couch (i.e., away from the gantry) and the needle end is identified. The coordinates of the ends of the needles comprise the input in the optimisation phase.

The parametric model involves the source-to-isocentre distance (SAD), the source-to-imager distance (SID), and a few distortion parameters describing the deviation of the machine from the 'ideal case'. Distortion parameters include the lateral and longitudinal imager offsets, distortion angles, and the coordinates of the central needle in the isocentre coordinate system. The parameter

values are obtained via optimisation (standard χ^2 fit) using the MINUIT routines of the CERN library [1]. The residuals in the lateral and longitudinal directions (versus the gantry angle) are obtained next; being zero in a perfect world, they reflect the mechanical instability of the entire machine. Typical results are shown below in Figs. 5a and 5b.

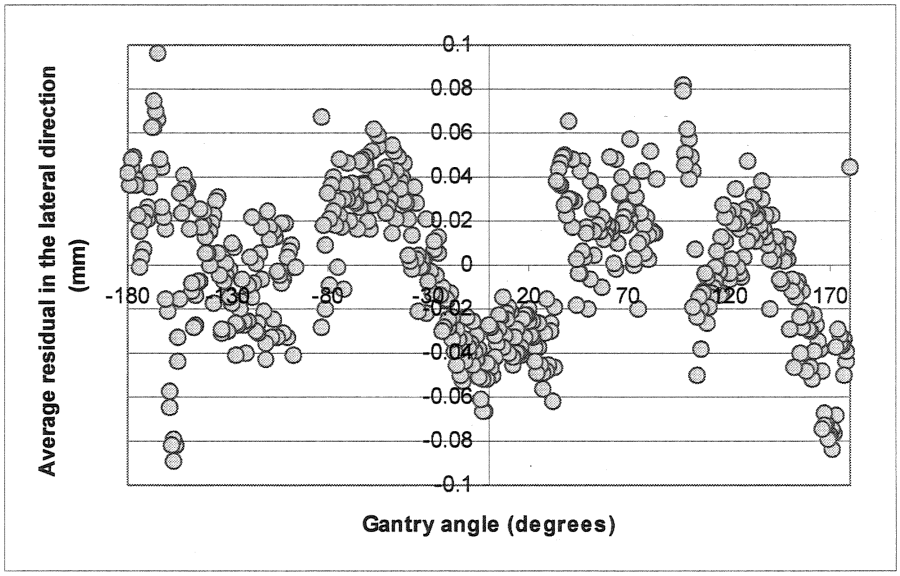

Figure 5a. Dependence of the lateral residual on the gantry angle

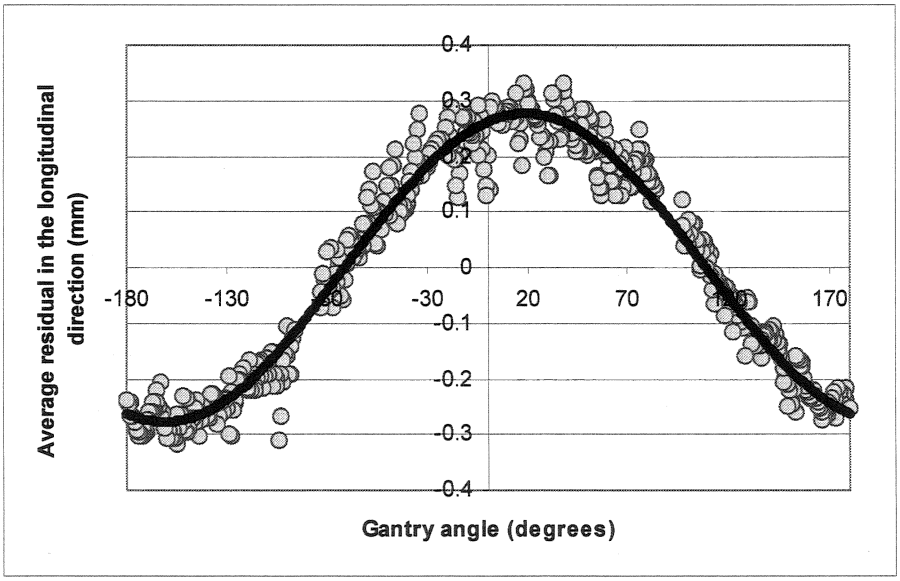

Figure 5b. Dependence of the longitudinal residual on the gantry angle

Evidently, the mechanical instability within one rotation of the gantry is small, namely, below 0.2mm in the lateral and 0.6mm in the longitudinal direction (full range). The movement in the longitudinal direction may be easily modelled on the basis of a cosine form with a phase shift; after this correction is applied, the mechanical instability in the two orthogonal directions is the same and amounts to about ±0.1mm.

The I_0 calibration

Assuming that the X-rays travel along straight lines, a CT image is essentially a map of the effective mass per unit surface (called T_s) along the line joining the source and the imager location where the ray is finally absorbed. To simplify the problem, a monoenergetic beam is assumed; the exact formalism may be found in Ref. [2]. Let I_0 and I_L be the incident- and transmitted-beam intensities along the ray line. T_s is then given by the formula:

$$T_s = \frac{1}{m}\ln\left(\frac{I_0}{I_L}\right) \qquad \text{(Equation 1)}$$

where m is the mass-attenuation coefficient corresponding to the energy used. Extracting T_s from I_L necessitates the knowledge of I_0. The I_0 map (i.e., the 'air counts') is produced by taking an exposure in open-field geometry (i.e., with air in the beam line).

The norm-phantom calibration

The norm-phantom calibration is essentially an I_0 calibration taken at a density different to that of the air. One may deduce that T_s is now given by the formula:

$$T_s = T_s^N + \frac{1}{m}\ln\left(\frac{I_L^N}{I_L}\right) \qquad \text{(Equation 2)}$$

T_s^N and I_L^N are, respectively, the effective mass per unit surface and the transmitted-beam intensity in the case of the norm phantom. The advantage, using a norm-phantom (instead of air) calibration, is that the second term on the right-hand side of Eqn. 2 is small in case that the density of the norm-phantom is close to the density of the actual object that is scanned; this is expected to lead to the reduction of non-linear effects (associated with the Hounsfield-units calibration).

The calibration for the beam-hardening correction

According to Eqn. 1, there is a linear correlation between the logarithm of I_0/I_L and T_s; the slope is equal to the mass-attenuation coefficient for the beam energy used. Physical beams, however, are polychromatic; this implies the existence of a spectrum of mass-attenuation coefficients, associated with the scanned object, not just one single number. Taking into account that the mass-attenuation coefficient is a decreasing function of the beam energy (for the energies used in the CBCT project),

it is expected that, in relative terms, the transmitted beam be richer (than the incident one) in high-energy components; this is because the low-energy components experience higher attenuation. The calibration for the beam-hardening correction uses a large cylindrical phantom and leads to the determination of the (first-order) correction terms in Eqns. 1 and 2.

The calibration of the dual-gain mode

Our FPDs may be operated in a number of output modes; a mode is identified as a set of options pertaining to binning, gain choice, etc. At present, our preferences lie with the dual-gain 2X2 mode[3]; in this mode, the charge, trapped in the semiconductor, is 'read out' by using (successively) two different capacitors, thus, leading to the so-called high- and low-gain signals. In this mode, each scan (approximately 650 projections) produces about 2GB of data that has to be transferred (from the acquisition system) to the workstation for additional processing (correction, storage, and reconstruction).

The dual-gain mode is an interesting option because of the following reasons. Quality in the reconstructed image is related, among other issues, to the noise level in the individual images. There are several sources of noise in the data, some of them being random (e.g., the quantum noise), others being systematic. As quantum noise is expected to be lesser of an issue (in relative terms) with increasing signals, one is tempted to increase the delivered dose. This strategy, however, solves one problem at the expense of creating another, inasmuch as signal saturation will then be observed in the low-attenuation regions of the same object; the saturation is due to limitations in the analogue-to-digital conversion of the signal[4]. Optimally, one could take two images of the object, corresponding to different dose levels, and combined them appropriately in order to achieve reasonable signals all over the projected image. By processing this 'dual' signal, one essentially achieves the enlargement of the dynamic range of the detector. Despite the fact that there is one value of the dose delivered in the case of the dual-gain mode, the way that the panel is read out creates a dual signal; one then uses the high-gain signal in the high-attenuation areas of the irradiated object and the low-gain one (properly scaled up) in the areas where the high-gain signal saturates.

The response of each individual pixel of the imager to the intensity of the radiation is carefully investigated. For each pixel, a set of parameter values is produced, comprising high- and low-gain dark fields, slopes (on the signal-versus-intensity plot), and a threshold value (maximal signal fulfilling the linearity in the signal-versus-intensity response) that determines which of the two signals is to be used. Fig. 6 shows one typical example of the evaluation for one pixel chosen at random.

---

[3] In order to minimise the processing time while still ensuring high-quality images, VMS is currently developing a dynamic-switching mode.

[4] At present, the output of our detectors is 14-bit limited.

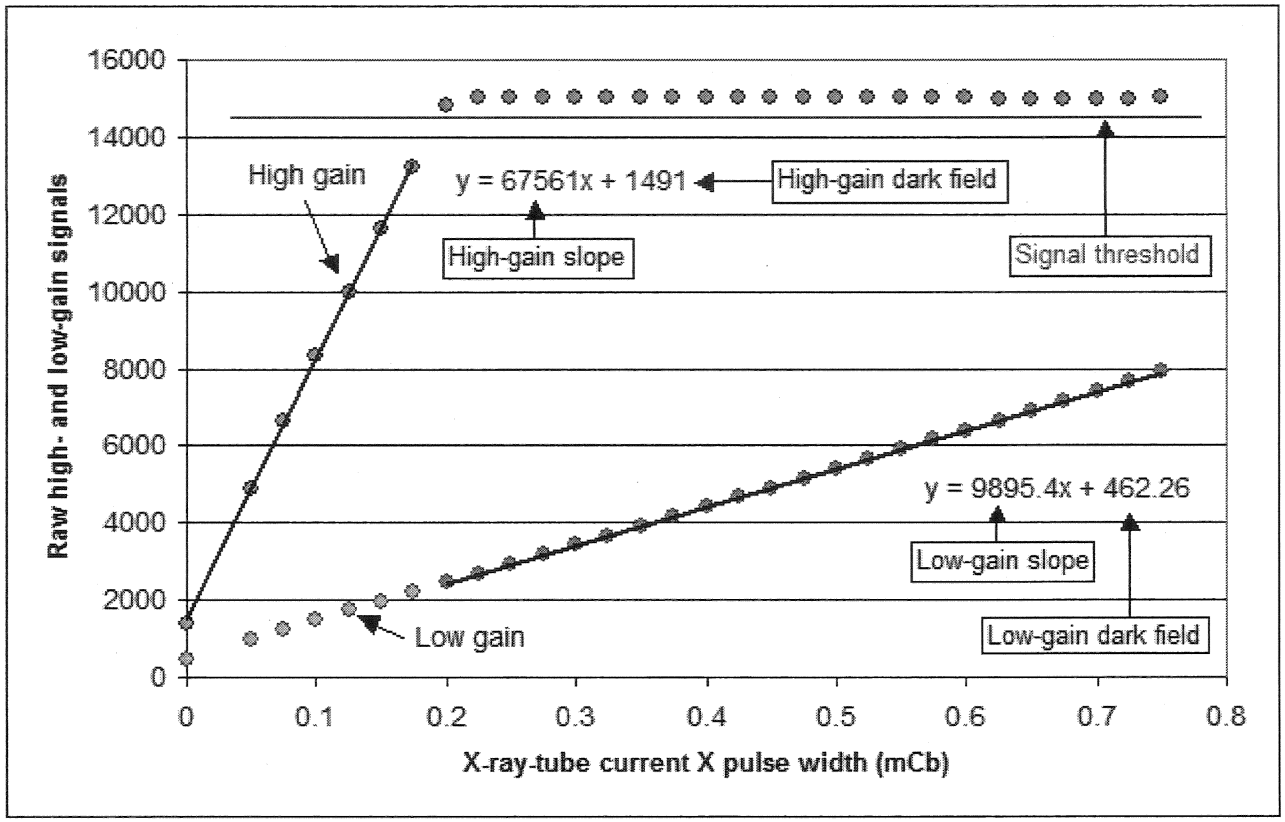

Figure 6. A typical example of the variation of the high- and low-gain signals with increasing intensity

All these values are finally stored and are subsequently used in the processing of the actual images. Defect pixels are interpolated.

The effort that has been invested in the calibration of the dual-gain mode is reflected neither in the simplicity nor in the length of this chapter. To mention an example, one of our challenging problems involved the ostensible non-linear response of the imager to the intensity of the incident radiation. Eventually, it was found that the culprit was not our FPD, but the pulse generator; the widths of the produced pulses did not agree well with the actual acquisition settings. We also found out that the (relative) deviation from the 'expected' value depended on the actual pulse-width setting. Small pulse widths (say, between 2 and 5msec) were affected more; the correction for the 2msec pulse may be as high as 20%. A scheme was carried out deriving the pulse-width corrections from the low-gain signal. For each pixel of the imager, two low-gain values were assumed 'safe' to use, namely, the measured dark field and the signal at the highest intensity used in the calibration (i.e., 30msec). The amount that each intermediate pulse-width value had to be corrected in order to bring the corresponding entry onto the straight line (defined by the two extreme measurements) was estimated. An average value (over the entire imager area) of the correction was calculated for each pulse-width value separately. The correction factors, thus obtained, were directly applied to the high-gain signal and were found to account beautifully for the non-linear effects that had been previously observed. It was relieving to realise that these corrections die off at pulse-width values

just short of 20msec.

Example of a reconstruction

Fig. 7 shows the reconstructed central slice of a head phantom using the present state-of-the-art knowledge; the dual-gain 2X2 mode was used in the data acquisition, the X-ray-tube voltage was 125kVp, and a 0.5mm-thick copper filter had been used for beam filtration. The reconstruction resolution was 512X512. The reconstruction has been done using the weighted filtered backprojection algorithm [3] combined with modified-Blackman noise filter. The input comprises 555 dual-gain images taken within one complete rotation of the gantry.

Figure 7. Reconstructed central slice of a head phantom

Future plans

To continue improving the image quality in the CBCT project, VMS is currently investigating a number of interesting effects. Concerning the imager development, we are mainly concentrating our efforts in suppressing lag and ghosting effects. On the side of the reconstruction, the elimination of ring and streak artefacts is our first priority, followed by the reconstruction problem in case of bones or implants. The issue of subtracting the scattering contribution from the images will soon be addressed.

To expedite the data acquisition and processing, VMS is about to release the so-called dynamic-switching mode; among others, this mode will give us the opportunity to increase the gantry-rotation speed and, consequently, to reduce the scan time, perhaps down to about 20sec. VMS is used to set the standards as high as possible and, keeping this in mind, we foresee that, in the near future, the whole process of CT-data acquisition, correction, transfer, and reconstruction will not last longer than about 2min.

Conclusions

With image-guided radiotherapy (IGRT), tracking and targeting the tumours will become simultaneous processes. This is needed in order to improve the efficiency, reliability, and safety in radiotherapy treatment.

Varian Medical Systems (VMS) has already produced and installed the first IGRT machine; the device comprises the VMS Clinac equipped with the On-Board Imager (OBI) component. Cone-beam CT (CBCT) imaging, one of the options of the OBI machine, aims at high-quality volumetric reconstruction of the irradiated object. The information on the actual positioning of the tumour and of the vital organs and tissue is processed quickly and modifications in the treatment plan are currently made on a daily basis. This enables the delivery of higher dose in each treatment session, is safe and accurate, and spares the healthy tissue surrounding the tumour. The dynamic targeting, i.e., the simultaneous tracking and targeting of tumours, is our next goal.

A number of calibrations are needed in order to operate a CT-imaging machine properly; these calibrations ensure that the machine components are properly aligned, the mechanical distortions are small, and yield important output that is used in the reconstruction of the actual scan data. The geometrical calibration is achieved by using a needle phantom placed close to the isocentre with the help of a laser system.

To boost image quality in the CBCT project, VMS has developed the dual-gain mode, a way to obtain useful data simultaneously in the high- and the low-attenuation areas of the irradiated object. In the near future, we will release the dynamic-switching mode, which is expected to do what the dual-gain mode is currently doing, but faster.

An indispensable part in our research program lies with creating the (software) tools to eliminate (ring, streak, and beam-hardening) artefacts in our reconstructed images. The detector development (to suppress lag and ghosting effects) is also high on our agenda. Finally, part of our activity pertains to the reconstruction problem in the presence of dense material, e.g., of bone or implants.

Acknowledges

The author acknowledges useful discussions with Drs. R. Suri and W. Kaissl. Dr. R. Suri presented the material, contained herein, at the SPIE International Symposium on Medical Imaging, San Diego, California, USA, February 12-17, 2005.

References

[1] James F 1996 MINUIT Minimization package *CERN Program Library Long Writeup D506*

[2] Joseph P M and Spital R D 1978 A method for correcting bone induced artifacts in computed tomography scanners *Journal of Computer Assisted Tomography* **2** 100-108

[3] Kak A C and Slaney M 1988 *Principles of Computerized Tomographic Imaging* (New York: IEEE Press)