Medical Image Analysis (1999) volume 3, number 3, pp 301–319

Computer-integrated revision total hip replacement surgery: concept and preliminary results

Russell H. Taylor[1]*, Leo Joskowicz[2], Bill Williamson[3], André Guéziec[4], Alan Kalvin[4], Peter Kazanzides[3], Robert Van Vorhis[3], Jianhua Yao[1], Rajesh Kumar[1], Andrew Bzostek[1], Alind Sahay[1], Martin Börner[5] and Armin Lahmer[5]

[1]Computer Science Department, The Johns Hopkins University, Baltimore, MD, USA
[2]Institute of Computer Science, The Hebrew University, Jerusalem, Israel
[3]Integrated Surgical Systems, Davis, CA, USA
[4]IBM T. J. Watson Research Center, Yorktown Heights, NY, USA
[5]Berufsgenossenschaftliche Unfallklinik, Frankfurt am Main, Germany

Abstract
This paper describes an ongoing project to develop a computer-integrated system to assist surgeons in revision total hip replacement (RTHR) surgery. In RTHR surgery, a failing orthopedic hip implant, typically cemented, is replaced with a new one by removing the old implant, removing the cement and fitting a new implant into an enlarged canal broached in the femur. RTHR surgery is a difficult procedure fraught with technical challenges and a high incidence of complications. The goals of the computer-based system are the significant reduction of cement removal labor and time, the elimination of cortical wall penetration and femur fracture, the improved positioning and fit of the new implant resulting from precise, high-quality canal milling and the reduction of bone sacrificed to fit the new implant. Our starting points are the ROBODOC$^{(R)}$ system for primary hip replacement surgery and the manual RTHR surgical protocol. We first discuss the main difficulties of computer-integrated RTHR surgery and identify key issues and possible solutions. We then describe possible system architectures and protocols for preoperative planning and intraoperative execution. We present a summary of methods and preliminary results in CT image metal artifact removal, interactive cement cut-volume definition and cement machining, anatomy-based registration using fluoroscopic X-ray images and clinical trials using an extended RTHR version of ROBODOC. We conclude with a summary of lessons learned and a discussion of current and future work.

Keywords: computer-assisted surgery, image-based registration, orthopedics, revision total hip replacement, ROBODOC$^{(R)}$ and ORTHODOC$^{(R)}$ systems

Received June 18, 1998; revised November 6, 1998; accepted November 10, 1998

1. INTRODUCTION

This paper describes an ongoing project to develop a computer-integrated system to assist surgeons in revision total hip replacement (RTHR) surgery (Joskowicz *et al.*, 1995). In RTHR surgery, a failing orthopedic hip implant, typically cemented, is replaced with a new one by removing the old implant, removing the cement and fitting a new implant into an enlarged canal broached in the femur. As the installed base of orthopedic implants grows and ages, replacement of existing implants, especially those relying on bone cement for fixation and fit, is steadily increasing. In 1992, 23 000 RTHR procedures were performed in the USA, and 27 000 in 1997. The average cost per procedure was $23 774 with an average hospital stay of 10.9 days (Orthopedic Network News, 1992).

*Corresponding author
(e-mail: rht@cs.jhu.edu)

RTHR surgery is a difficult procedure fraught with technical challenges and a high incidence of complications. Femoral cement removal and canal preparation present the most difficulties (Turner and Emerson, 1982; Crenshaw, 1987). The goal is to remove as much of the old cement as possible to facilitate the insertion of a new implant and to provide an optimal surface for bone support and interdigitation. While the cement mantle in the proximal area of the canal is visible and easily accessible, the cement mantle and plug in the distal area are hard to see and reach due to the canal depth and the bowing of the femur. Removing cement is tedious, time consuming and risky, taking on average between 30 min and 2 h. Femoral canal preparation for a revision is more difficult because there is less good bone left and because the surgical manipulations are more delicate. The reamers tend to follow the old canal, making axis and canal position corrections virtually impossible. The femur is fractured in about 18% of cases, and the surgeon breaks through the cortical wall of the femur in another 10% of cases (Schurman and Maloney, 1992). When errors occur, more time is required to repair the damage, additional blood is lost and the risk of infection increases.

None of the current techniques for cement removal are fully satisfactory. Osteotomes and flexible reamers are difficult to manipulate and have the tendency to follow the pathway of the old canal. Hand-held high-speed drills cut cement fragments but require fluoroscopic X-ray images obtained with an intraoperative C-arm for careful guidance to avoid perforating the femur walls. Fluoroscopic guidance requires repeated imaging and is error-prone because it provides static, uncorrelated two-dimensional (2-D) images and results in cumulative radiation exposure to the surgeon. Lateral femoral windows facilitate distal access to the cement but may compromise bone integrity. A recently developed method uses low-viscosity cement that bonds to the old cement to form a plug. The plug is then pulled out by screwing in a threaded extraction rod and pulling out pieces in short segments. This technique cannot be used when the cement mantle widens distally or when cavities are present in the side of the bone. New cement removal technologies, such as cement softening with ultrasonically driven tools or cement fracturing with a lithotripster can lower the complication rate but are unlikely to significantly improve accuracy or shorten the procedure.

The growing numbers, greater difficulty and reduced margin for error make RTHR surgery a natural target for robotic machining. Our goals in developing the system are:

(i) the significant reduction of cement removal complications, specifically cortical wall penetration and bone fracture;

(ii) the significant reduction of cement removal labor and time required;

(iii) improved positioning accuracy and fit of the new implant resulting from precise, high-quality canal milling;

(iv) the reduction of bone sacrificed to fit the new implant; and

(v) the reduction of cumulative surgeon exposure to X-ray radiation caused by the repeated use of the fluoroscopic C-arm.

In addition to the direct patient benefits, these advantages can save costs, both by reducing operating room cost and by shortening the hospital stay and recovery time.

In 1995, we initiated a joint effort to develop an integrated preoperative planning and intraoperative execution system for RTHR surgery. Our starting points were ROBODOC$^{(R)}$ (Paul *et al.*, 1992; Taylor *et al.*, 1994), Integrated Surgical Systems' (ISS) computer-integrated system for primary total hip replacement procedures and the manual RTHR surgical protocol. ROBODOC was developed clinically by ISS from a prototype developed at IBM Research and is currently in clinical use in Europe. Preclinical testing showed an order-of-magnitude improvement in precision and repeatability in preparing the implant cavity. Over 3000 human cases have been performed to date (fall of 1998), with very positive results documented in follow-up studies (Börner *et al.*, 1998). In primary total hip replacement (PTHR) procedures, the damaged joint connecting the hip and the femur is replaced by a metallic implant inserted into a canal broached in the femur. The ROBODOC system allows the surgeons to plan the procedure preoperatively by selecting and positioning an implant with respect to a computer tomography (CT) study and intraoperatively mill the corresponding canal in the femur with a high-speed tool controlled by a robotic arm.

The ROBODOC system consists of an interactive preoperative planning system, called ORTHODOC$^{(R)}$ and a robotic system for intraoperative execution. ROBODOC PTHR surgery starts with a minor surgical procedure in which two small pins are implanted in the distal and proximal femur surface (earlier, a third distal pin was required). A CT scan of the patient, showing the femur and the implanted pins, is then acquired. Next, ORTHODOC processes the CT dataset, locates the pins within the CT images, constructs a three-dimensional (3-D) surface model of the bone, and shows the surgeon three orthogonal planar slices through it. The surgeon selects a desired implant model and size and interactively positions a CAD model of it with a mouse relative to the slices. ORTHODOC generates cross-sectional displays of the implant model showing the planned placement

superimposed upon the planar sectional views selected by the surgeon. In the operating room, surgery follows the usual protocol up through the point where the femoral head is removed. The femur is then placed into a fixation device attached to the robot's base. The pins are exposed and located in robot coordinates by a combination of force-compliant guiding and autonomous tactile search by the robot. The system then computes the transformation from CT (planning) to robot and patient (actual) coordinates and machines out the canal shape associated with the chosen implant while the surgeon follows the progress on a screen. Once the shape is cut, the robot is moved out of the way and the procedure is completed manually as usual.

RTHR surgery is more complex than PTHR surgery: it can have more uncertainty associated with it and requires more system capabilities. Surgeons must plan for and remove the old implant and cement before cutting the new canal cavity. They must plan for the new cavity in the presence of the old implant and cement with lower-quality CT images containing imaging artifacts caused by the old metal implant. They must foresee complications in implant and cement removal, which might change or invalidate the preoperative plan. Consequently, computer-assisted RTHR surgery requires substantial extensions and modifications to the ROBODOC PTHR surgery paradigm. To summarize, the system must provide, in addition to the current capabilities, cement removal planning and cutting, intraoperative plan modification and uncertainty assessment, and the integration of intraoperative fluoroscopic X-ray images with preoperative CT data.

In this paper, we discuss the main difficulties of computer-integrated RTHR, identify the key technical challenges, and investigate possible solutions. We propose a new system and protocol for preoperative planning and intraoperative execution. We present a summary of methods and preliminary results in CT image metal artifact removal, interactive cement cut-volume definition and cement machining, anatomy-based registration using fluoroscopic X-ray images and clinical trials using an extended RTHR version of ROBODOC. To obtain accurate anatomy-based registration, we describe methods and experiments for fluoroscopic X-ray image distortion correction, C-arm camera calibration, and X-ray and CT bone contour and surface extraction. We conclude with a summary of lessons learned and discuss current and future work.

2. PROBLEM ASSESSMENT AND REQUIREMENTS

To identify the requirements of computer-integrated RTHR surgery and gain an understanding of the practical problems, we follow the steps of the manual RTHR procedure with the ROBODOC PTHR protocol. We identify the differences, missing components and assess the adequacy of the current techniques. We evaluate the relative importance of the difficulties that arise and explore systematically possible solutions to them.

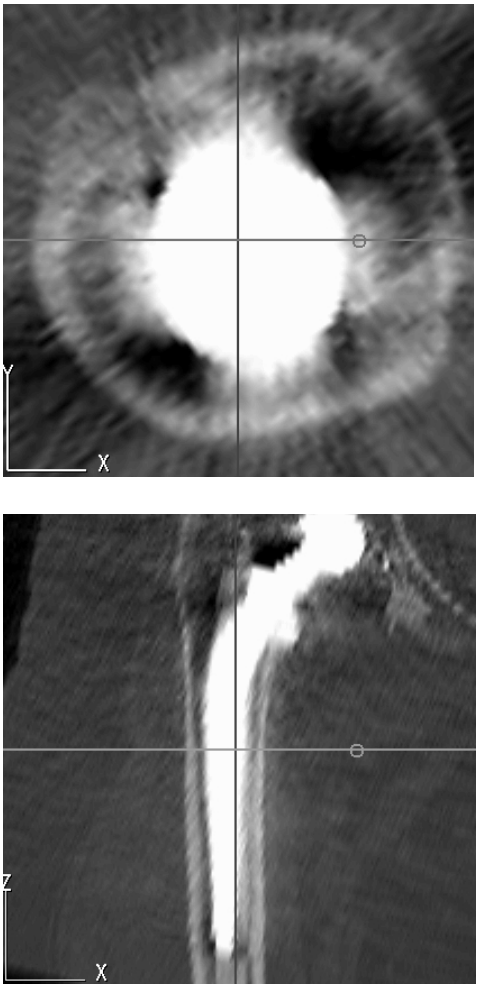

Figure 1. Original cross-sectional slice of a failing implant (top) and ORTHODOC reconstructed frontal view (bottom).

2.1. CT images

X-ray CT images of body sections containing metal objects are frequently corrupted by reconstruction artifacts that often resemble streaks radiating from the regions of the image where metal is present (Figure 1). Because metal objects are opaque to X-ray beams in the diagnostic energy range, their scanning yields incomplete projection data. CT images

reconstructed from this incomplete data contain imaging artifacts, whose extent depends on the implant material and volume. Imaging artifacts in CT datasets of RTHR patients with metal femoral implants are most marked in the proximal section, where the implant is the thickest. The imaging artifacts make it difficult to determine the boundary between the implant, the cement and the bone. Since the quality and accuracy of the surgical plan depends directly on the quality of the CT images, reducing image artifacts as much as possible is of great importance.

2.2. Preoperative planning

Preoperative planning of RTHR surgery consists of cement removal and new implant planning. Cement removal planning defines the cut volume that contains as much of the old cement as possible. New implant planning determines the type, size and position of the new implant and the associated canal cut volume that guarantees a precise fit. Both steps are interrelated, since the bone stock left after cement removal determines the implant types, sizes and positions that can be used. Conversely, the available implant types and sizes determine the new canal shapes, which indicate what bone and cement volumes should be removed and what contacts and gaps will appear once the new implant is in place.

The main difficulties of RTHR preoperative planning stem from the uncertainties associated with imaging artifacts, old implant removal and cement removal. Specifically, four issues must be addressed:

(i) determining the extent of the cement mantle and the bone stock requires substantial experience and judgment from the surgeon. Cement tends to partially fill porous bone, creating heterogeneous zones that must be evaluated individually;

(ii) the cement cut volume must be defined individually for each case. A fast, intuitive accurate method must be developed to allow surgeons to define machinable cut volumes from CT data;

(iii) determining the trade-offs between the cement cut volume and the canal cut volume. The cement cut shape must be compared with the implant cut shape to identify discrepancies, such as cement pockets and lack of surface support for the new implant. Because there are no generally agreed upon criteria for the best strategy, automatic cement cut-volume determination is unlikely;

(iv) determining the shape and extent of the cement mantle that will be left after the old implant is removed. When the old implant is not loose, bone detachment and even femoral fracture can occur during implant removal, thus invalidating the preoperative plan. The options are to intraoperatively modify the plan, to create alternative backup plans or to complete the procedure manually.

2.3. Intraoperative validation and re-planning

To account for the uncertainties introduced by the old implant and cement removal, preoperative plans must be compared and validated against the intraoperative situation. This validation is not necessary for the PTHR ROBODOC procedure since the femoral anatomy does not change before the canal is cut. It might be necessary for certain RTHR cases, when the surgeon needs to gain confidence in the preoperative plan and possibly to modify it with additional intraoperative information. The modifications include changing the cement cut volume and shape to allow for more or less cement removal and adjustments to the new implant size and position.

The key difficulties are the integration of the intraoperative data with the preoperative plan and the modification of the preoperative plan. Currently, the sources of intraoperative data are visual and tactile inspection and fluoroscopic X-ray images of the canal and remaining cement mantle. To be useful, this data must be integrated and correlated with the preoperative plan and CT study. Any modifications to the preoperative plan must be done quickly and accurately.

2.4. Image, patient and robot registration

Robotic procedures require registering preoperative plans and images to the robot and to the actual patient anatomy. The ROBODOC system for PTHR surgery uses a proximal and a distal pin implanted prior to surgery in the femur before the patient is scanned. The robot is registered to the patient by bringing its tip in contact with the fiducials. The robot and the preoperative plan are registered by locating the pin positions on the CT images. In RTHR surgeries, osteotomy of the greater tronchanter is often necessary to provide better exposure and ease the insertion of the new implant. Thus, a new location that does not interfere with the cutting tool or require more invasive surgery must be found. Alternative methods, such as the use of external fiducials, the acquisition of points on the bone surface or the use of fluoroscopic images for anatomy-based registration require investigation to determine the most practical and robust solution.

2.5. Cement and bone cutting

Once the femur has been fixated and the robotic arm has been registered, the cement and the canal are machined. A revision canal contains cement and is one-third to one-half longer than a primary canal, extending below the bow of the isthmus. Because of its extended length and curvature, machining the new canal raises issues of robot reachability and workspace capabilities, the stiffness of the robotic arm

and cutting tool, and accessibility for curved machining paths. Experiments must be carried out to determine the effects of cutting cement with high-speed tools: stresses and femur fracture analysis, accuracy, rough versus finish cutting passes. These will establish whether the ROBODOC tool wrist, as well as the cutting strategy must be redesigned.

3. SOLUTIONS ANALYSIS AND PREVIOUS WORK

We have identified four major issues that must be addressed to develop a practical RTHR surgery system:

(i) preoperative planning under uncertainty due to imaging artifacts, old implant removal and cement removal;

(ii) cement cut-volume definition and implant selection for each individual case;

(iii) intraoperative plan validation and modification; and

(iv) image, patient and robot registration.

3.1. CT image artifact reduction and preoperative planning

CT image metal artifact removal is a natural starting point for attempting to reduce the uncertainty associated with preoperative planning. Several approaches have been suggested, including:

(i) using implants made of materials with lower attenuation coefficients or with smaller cross-sectional areas;

(ii) using higher-energy X-ray beams that will not be blocked by the implants;

(iii) averaging out the effect of the imaging artifacts by interpolating and reslicing the 2-D images stack (Robertson *et al.*, 1988);

(iv) averaging out the effect of the imaging artifacts by combining multiple image sets, each scanned with the gantry at a different angle;

(v) interpolating the missing projection data and reconstructing the images from these completed projections (Oppenheim, 1977; Hinderling *et al.*, 1979; Glover and Pelc, 1981; Kalender *et al.*, 1987; Medoff, 1987; Klotz *et al.*, 1990); and

(vi) creating simulated projection data from the images, interpolating the missing data in these projections and then reconstructing the images (Tuy, 1992).

For reducing imaging artifacts in CT data of RTHR patients, (i) is obviously not an option. The need for limited patient dose and low energy to discriminate among materials (biological tissue types and synthetic material such as cement) rules out (ii). The averaging effect of (iii) reduces not only image artifacts but also image resolution, while (iv) requires longer scanning time and higher patient exposure to radiation. In principle (v) can produce the best results, but in practice access to raw projection data is problematic for a clinically viable solution. Option (vi) is the most practical and has the added advantage that any methods developed can readily be applied to real projection data when these are available. However, both (v) and (vi) are intrinsically limited because crucial image details are mostly erased by the imaging artifacts and cannot be recovered from the CT image data alone.

One way to improve the information available for preoperative planning is to complement the CT information with X-ray images such as digitized multi-planar films, fluoroscopic images or CT scout images. Since these images have no reconstruction artifacts, they can provide some of the missing contour information. To be useful, several images must be captured, accurately registered to the CT study and properly presented to the surgeon. It remains to be determined how well the surgeon can position the implant and define a cement cut volume with these two kinds of images. A novel alternative is to use scout images taken at the time of the CT study and use those to fill in the missing data. The CT scanner provides the precise data for correlating the scouts with the CT data. We describe this approach in Subsection 5.1.

Another possibility is to use several co-registered X-ray images instead of a CT study (in fact, a crude version of this method is currently used to plan manual surgeries with acetate overlays). The disadvantage is that much less volumetric data is available for planning, although it is inexpensive. A key issue is the number of images necessary and how to co-register them.

3.2. Cement cut-volume definition

A custom cement cut volume must be created for each patient. The creation process must be fast, intuitive and produce an accurate, machinable shape. Cut volumes can be created by specifying, for each CT slice, points defining a 2-D contour bounded by splines. The stack of 2-D slices defines a 3-D cut volume. Adjusting the cut-volume shape to fit the cement requires moving the points. While accurate, this method requires the surgeon to input and manipulate many points. An alternative is to define the contours in a subset of the slices and interpolate the rest automatically. This trades off the number of input points for accuracy of the matching shape.

Another option is to have one or more simple, parameterized shapes, such as cones with elliptical cross-sections, and fit them to the cement by varying the parameters. Although modifying a few parameters is fast, this method is potentially unintuitive and might take longer to produce the desired results. A hybrid method which uses simple parameterized shapes for the rough fit and control point modification for fine tuning could provide the best trade-off.

The cut-volume shape thus defined must then be approximated to a machinable shape determined by the radius of the cutter and the machining axis. The smaller the cutter radius, the more accurate the shape, but the longer it takes to machine. Multi-axis machining methods, in which the cutter axis is repositioned during cutting, allow tighter fits but require more complex computation. Once the cut-volume shape has been defined, the new implant and its associated canal must be selected and positioned. The implant can be selected and positioned manually, as in the current ORTHODOC system or by interactive definition of implant and image landmark correspondences.

A comparison of the cement cut shape and the positioned implant cut shape is necessary to identify discrepancies, such as pockets and lack of surface support for the new implant. The comparison can be left to the surgeon, by overlaying the two volumes graphically and showing them in different views. Reconciling discrepancies between the two cut volumes can be difficult. An alternative strategy is to define a single cut volume for both the cement cut volume and the new canal cut shape by taking their union. In this scheme, the new implant size and position is chosen so as to contain all or most of the cement and the old canal. The cut volume associated with the new implant is then used to mill the old cement mantle, cement plug and new canal shape simultaneously. The advantage of this approach is that no new cut shape needs to be defined or modified. The disadvantage is that a trade-off must be made between removing all the old cement and removing too much good bone. Lumping old cement removal and new canal preparation assumes that the preoperative plan is of high quality, since no intraoperative adjustment is possible once the robot starts cutting the shape. Also, it diverges from current practice, which views cement removal and canal preparation as two distinct steps. This was found to be acceptable in some cases, as discussed in Subsection 5.8.

3.3. Intraoperative plan validation and modification

Fluoroscopic X-ray images provide the currently accepted and most practical way for intraoperative plan validation. Visual and tactile inspection rely on the surgeon's ability to mentally correlate the CT data to the intraoperative situation. This correlation is qualitative, and is only effective for detecting major discrepancies. Fluoroscopic X-ray images provide more accurate information than visual or tactile inspection but must be corrected for distortion and registered with the CT data to be useful in a robotic procedure (Yaniv *et al.*, 1998). Other intraoperative imaging techniques, such as CT or ultrasound devices, are either not yet available or have not proven to be sufficiently accurate.

The preoperative plan can be validated by superimposing the cut volume and the new implant projections on intraoperative fluoroscopic X-ray images. The surgeon can judge their adequacy visually and modify the plan accordingly. One approach is to choose the best of several preoperative plans. Another is to have the surgeon change the shape and size of the cut volumes interactively, as in the preoperative planning phase, although this can be impractical in the operating room. It is best to first determine the extent of the discrepancies and the nature of the modifications required before committing to a specific solution.

3.4. Image, patient and robot registration

Accurate robot-to-patient and robot-to-image registration are essential for RTHR surgery. They can be achieved by several methods, which include:

(i) implanted fiducials, as in ROBODOC;

(ii) acquisition of points on the bone surface by direct contact of the robot tip (the 'cloud of points' approach in Lavallée (1995); or

(iii) contactless registration using fluoroscopic X-ray images.

Methods based on implanted fiducials are reliable and accurate, but require an extra procedure. For RTHR surgery, a new position of the pins must be found. Cloud-of-point methods are only practical when enough bone surface is exposed to allow direct contact[a]. Pinless registration of X-ray images to CT data attempts to use correspondences between anatomical structure features in both images to achieve the match. This type of 2-D to 3-D anatomy-based registration has been discussed in Lavallée *et al.* (1994) and Lavallée (1995), and more recently incorporated in a robotic system by Brandt *et al.* (1997). It presents significant technical challenges but has clear advantages over the other methods.

The main technical issues for efficient, robust and accurate real-time anatomy-based registration of CT data with fluoroscopic X-ray images are:

(i) fluoroscopic X-ray images present geometric distortion of up to several millimeters that must be corrected;

[a]Integrated Surgical Systems has recently demonstrated a variation of this method for PTHR in the DigiMatch Single Surgery System. The system uses a passive mechanical digitizing arm to collect the cloud of points.

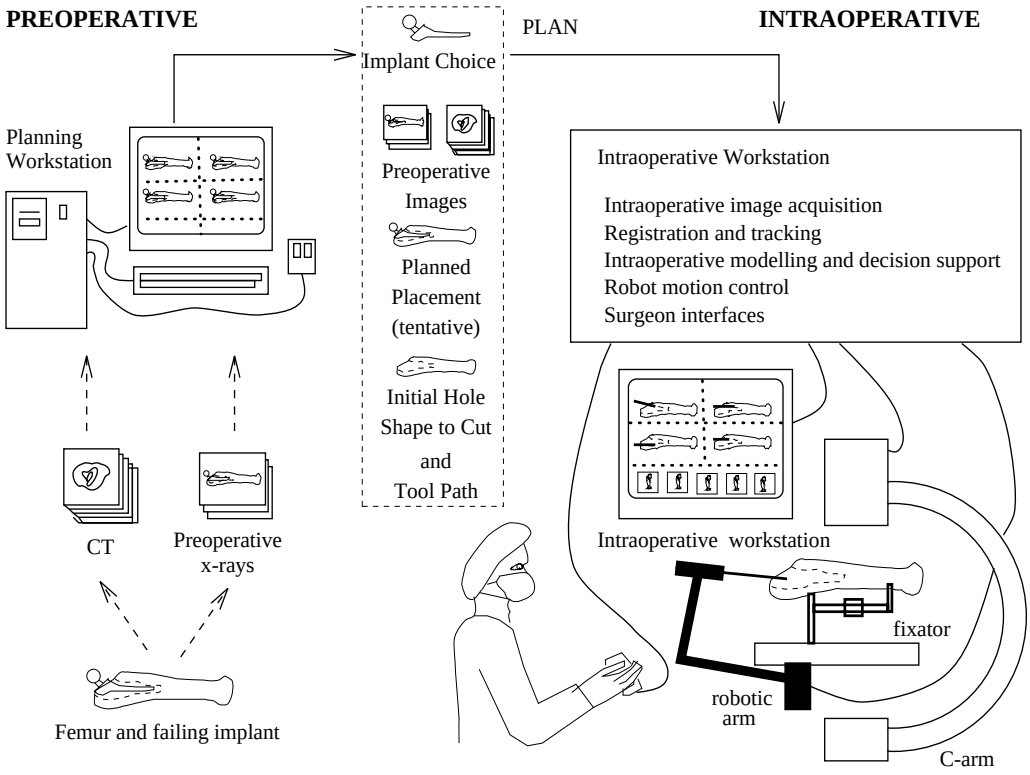

Figure 2. System for RTHR surgery.

(ii) the C-arm intrinsic and extrinsic parameters must be obtained for each dataset;

(iii) bone surface contours must be extracted from the CT dataset and the fluoroscopic X-ray images; and

(iv) the extracted contours must be matched to obtain the registration by minimizing the distance between the corresponding contour features.

Related work in this area includes (Brown, 1996; Joskowicz *et al.*, 1995; Lavallée, 1995; Hofstetter *et al.*, 1997; Viant *et al.*, 1997; Brack *et al.*, 1998; Hamadeh *et al.*, 1998). We present algorithms for accurate image dewarping, C-arm camera calibration, contour extraction and 2-D/3-D registration in Section 5.

4. PROPOSED SYSTEM ARCHITECTURE AND PROTOCOL

Figures 2 and 3 show the proposed experimental system architecture and protocol for computer-integrated RTHR surgery. Both are designed to accommodate a variety of technical solutions for the different modules and steps of the procedure. At any point in time, the actual system will consist of specific instances of modules and steps, which are then tested and refined based on experimental results and surgeon feedback.

Preoperatively, a CT scan and/or X-ray images of the femur with the failing implant are acquired. Fiducials will be first implanted when the robot/image registration is fiducial based. External fiducials are optionally attached to the table or to the patient to register the CT and X-ray datasets. If the X-ray images are CT scouts, there is no need for external fiducials or registration as the patient is kept still and the gantry of the CT machine is controlled precisely. The CT slices, together with the X-ray images and their view pose information are then loaded into the preoperative workstation. The images are then registered if necessary. The system then processes the CT images and additional X-ray images to reduce imaging metal artifacts.

For preoperative planning, the system presents enhanced CT images, or alternatively, when only X-ray images are available, an image spreadsheet which keeps the images co-registered and allows manipulating overlays on them is used. In either case, the images are used to evaluate the extent of the cement mantle, define a cut volume and select an implant

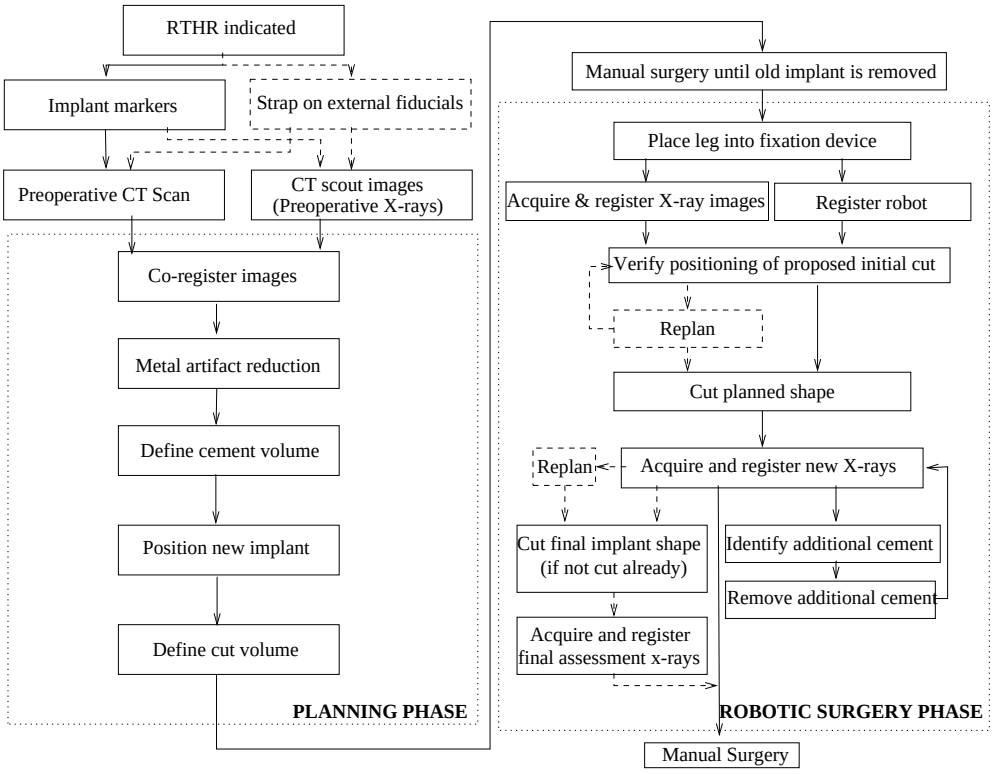

Figure 3. Procedural flow for RTHR Surgery. Dotted boxes and lines indicate optional steps and/or extensions.

type, size and position. The surgeon defines interactively the cut cement volume and implant type and placement, and the system produces a cut volume. The output is a set of co-registered preoperative images and one or more preoperative plans consisting of an implant type, size and position and the shape and position of the cement cut volume.

Intraoperatively, manual surgery will proceed as usual until the old implant is removed and the cement that can be easily removed manually is out. To remove the remaining cement, the surgeon will place the femur in a fixation device rigidly attached to the robot. The fixator must be designed so that it can hold fragile bones without damaging them. Registration between the patient and the robot will be achieved via the implanted pins or with intraoperative fluoroscopic X-ray images. If the preoperative plan needs to be validated, the system will display the outline of the volume to be cut superimposed on the intraoperative fluoroscopic X-ray images. The surgeon may adjust the surgical plan, either by repositioning the cut volume or by modifying its shape. To be practical, this process must take $<$10 min.

Once the surgeon has verified the plan, the designated volume will be cut out using the same material removal strategy as that employed by ROBODOC for PTHR surgery. Based on the surgeon's evaluation and confidence in the preoperative cut volume, the robot can then cut out either the entire volume or cut only a conservative initial volume. Additional images will be taken, registered and compared with the planned cut volume. These images will be used to assess what material still needs to be removed and to update the registration of the robot to the patient. The surgeon will then instruct the robot to remove additional cement volumes, optionally acquiring additional fluoroscopic X-ray images, and will select a final implant model and position using the most recent images. The robot will then cut the final shape. Once the femoral cavity for the new implant is prepared, the robot will be removed from the surgical field and manual surgery will proceed as usual.

5. PRELIMINARY RESULTS

We are currently implementing various steps of the procedure. The work has concentrated on CT imaging metal artifact removal, interactive cut-volume definition, anatomy-based registration using fluoroscopic X-ray images, image spreadsheets, cement machining experiments and clinical RTHR surgery experience with ROBODOC. To perform

anatomy-based registration to the desired accuracy, we developed methods and experiments for:

(i) fluoroscopic X-ray image distortion correction and C-arm camera calibration;

(ii) X-ray and CT bone contour and surface extraction; and

(iii) anatomy-based registration of CT and fluoroscopic X-ray images.

We describe each briefly next.

5.1. Imaging metal artifact removal in CT data

We have developed a new method for reducing imaging metal artifact noise in CT data using scout images. The scout images, which are acquired in the same session as the CT data, provide additional information necessary to reduce the imaging artifacts. Scout images are produced by the CT machine by keeping the X-ray tube in a fixed position and moving the table in constant increments. The main advantage of using scouts is that there is no need to register them to the CT slices. Also, the accuracy of the correlation between scout images is the same as that of CT slices and the residual error is negligible for our application. As with any CT study, the patient must lie still while being scanned.

In principle, it is possible to reduce imaging metal artifacts by first converting a dense set of scout images into a dense set of sinogram-like projections, then processing these projection data, and finally obtaining reconstructed CT images using a standard filtered back-projection method (Browne, 1988). However, our experiments show that to obtain adequate results, the number of scouts required is >100, which is clinically unacceptable both in terms of scanning time and radiation dosage. Instead, we have developed a method that combines the data from a small number of scout images with the noisy CT data. Our goals were to design an off-line method to recover object boundaries in the region around the metal implant, rather than recovering absolute pixel densities over the entire image.

We developed MARCUS (Kalvin and Williamson, 1998), an iterative algorithm that reduces the imaging metal artifacts in CT data. The algorithm repeatedly modifies the corrupted CT images by applying sets of image and projection constraints. Image constraints are based on information about the physical properties of the objects in the image, while projection constraints are imposed by scout data. MARCUS starts by updating the CT images according to the image constraints. Then, the modified images are checked for consistency with the scout data to reduce the inconsistencies between the two. This process is repeated until the change in the images in successive iterations is below a prespecified threshold, or until an iteration limit is reached. In addition, geometric bounds on the position of the implant contour in the CT images are computed by detecting the edges of the metal implant in the scout data and using the Radon transform basic property, whereby each point in the projection space corresponds to a line in image space.

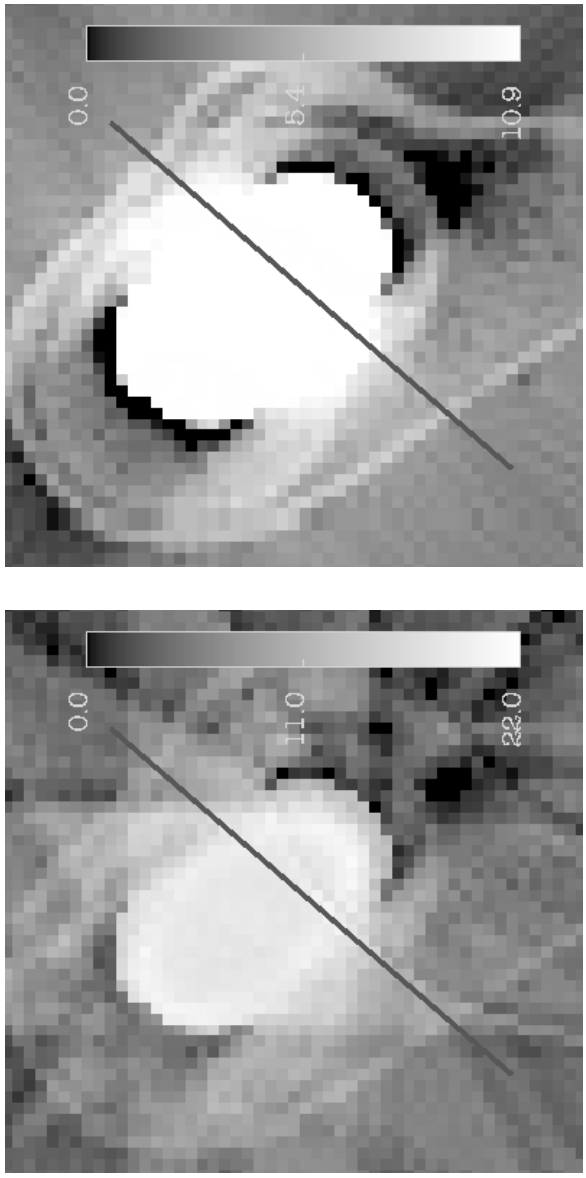

Figure 4. Detail of a CT image before (top) and after (bottom) 50 iterations of MARCUS imaging metal artifact reduction.

We have tested the MARCUS algorithm on a realistic custom-built RTHR phantom which consists of a cadaver

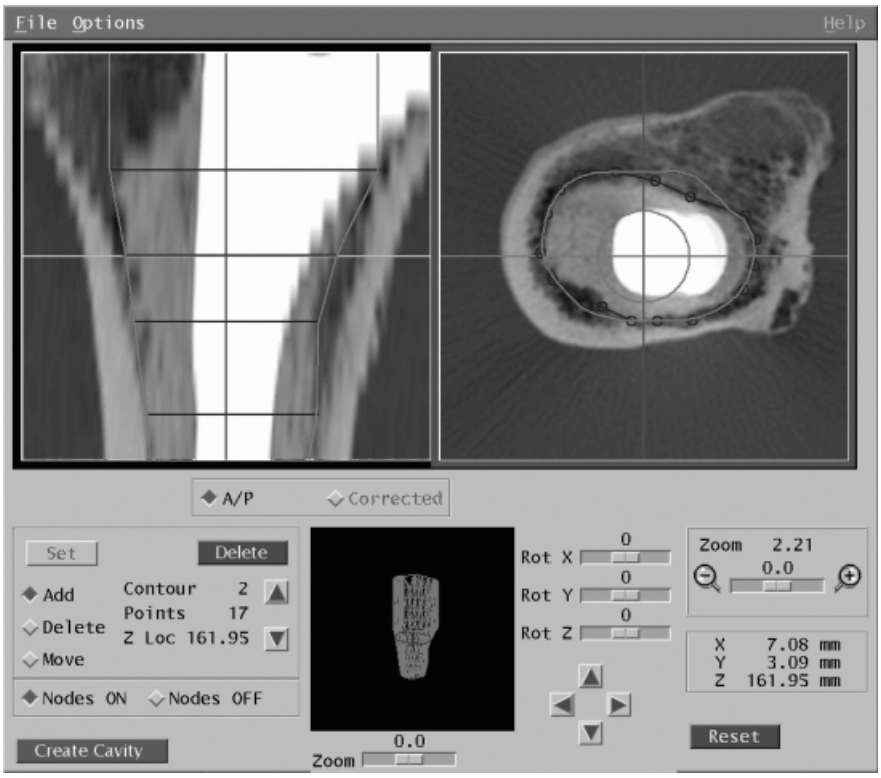

Figure 5. Extended ORTHODOC screen showing an A/P view (left) and a top view of interactive cut-volume contour definition (right). The desired volume is the exterior (outer line, white) and the computed (inner line, black) contour. The bottom screen shows a wire-frame 3-D view of the defined cut volume.

femur, a metal hip prosthesis and bone cement. Figure 4 shows a CT image before and after the reconstruction using 12 scout images. The effects of the imaging metal artifacts have been reduced significantly, and details of the implant boundary have been recovered. This experiment also confirms that the 'blooming' nature of imaging metal artifacts in CT data causes the complete loss of image information in the region close to the metal. This blooming is not simply a display artifact, e.g. resulting from saturation at the high end of the gray-scale map. We thus conclude that effective imaging metal artifact reduction cannot be achieved by processing the noisy CT images alone.

5.2. Interactive cut-volume definition

We have augmented ORTHODOC with an interactive cut-volume definition module. In the module, the surgeon first segments out the bone cement by creating a contour that defines the bone cement to be removed in several CT slices. The contour is created from control points supplied by the surgeon in each CT slice by first fitting a spline in each slice and then connecting the resulting contours in successive slices to form the volume. The splines can be modified by moving, adding and deleting control points. Currently, the volume is defined as a single connected piece, so the volume creation process is straightforward. Figure 5 shows the user-interface screen for interactive cut-volume definition.

The contour data are then input into a cut-path generator algorithm which outputs a contour identifying the computed robot cut path. The cut path is created by successively processing from top to bottom all the contours that the user has generated. The constructed single-axis machining cut path takes into account the cutter radius and allows for straight insertion along the vertical axis.

5.3. Fluoroscopic X-ray image distortion correction and C-arm camera calibration

We have developed several methods for performing intrinsic and extrinsic calibration of fluoroscopic X-ray images. Intrinsic imaging parameters include focal length, image center, pixel scaling and image warping, and can be computed by analyzing an image of a calibration object of known geometry. Extrinsic parameters describe the position and

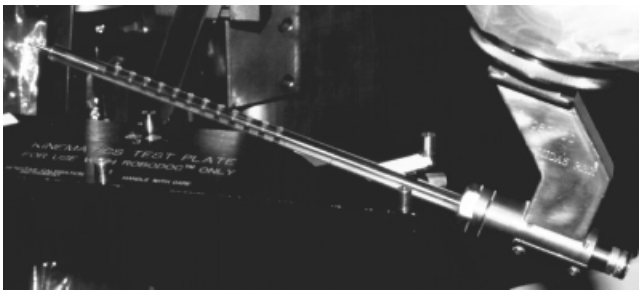

Figure 6. Radiolucent X-ray image to robot calibration rod.

orientation of the X-ray source with respect to an external reference frame such as the fixator or the femur. These parameters are used for co-registering multiple fluoroscopic X-ray images and for registering intraoperative X-ray images to preoperative–operative images and models. We are currently comparing them to determine which is most suitable for our particular intraoperative environment.

One method Guéziec *et al.* (1998) extends the NPBS method of Champleboux *et al.* (1992) to provide a direct registration of the robot to fluoroscopic X-ray images. Two significant differences are: (i) the use of thin-plate splines (Duchon, 1976; Bookstein, 1991) for dewarping interpolations and (ii) obtaining images of multiple known calibration points. Rather than having the robot hold a complex calibration object comprising one or more planar grids of radio-opaque balls (Champleboux *et al.*, 1992), we use a radiolucent probe placed into the cutter collet of the robot (Figure 6). The probe is swept and imaged through two planes to construct a virtual calibration grid. Image processing to recover the ball centers is simple and robust since the only difference between successive images is the position of the balls. The exposure radiation to acquire the multiple images is higher but still within acceptable limits. The advantage of this method is that it allows for more compact and convenient fixturing than would be feasible with a conventional two-plane grid. Our experiments show accurate 3-D localizations within 0.1–0.4 mm with a conventional C-arm (Guéziec *et al.*, 1998).

We have also investigated an alternative approach for intrinsic and extrinsic calibration on the same set of images by imaging a partially radiolucent aluminum plate mounted on the C-arm's image intensifier. The plate has 3/16 in thick and 1/16 in wide grooves machined in a square pattern on 1/4 in centers. The grooves appear as pale lines on the fluoroscopic X-ray images (Figure 7) and can be segmented with standard image processing techniques. The dewarping map is obtained by locating the image points of each grove's centerline, fitting a fifth-order Bernstein polynomial to them, and computing the dewarped image point positions with a scan line algorithm (Wolberg, 1990). The scan line algorithm computes the intersections between each horizontal scan centerline line in the image and the vertical Bernstein-based curves. The intersection points on the scan line are then used to fit a piecewise cubic spline by relating their location on the scan line against their ideal locations. The resampled image in the horizontal direction is obtained by interpolating all the pixels on the scan line using the cubic spline line. The same process is used to dewarp the grid in the vertical direction. Figure 7 illustrates the process. The advantages of our method over methods which use a grid of radio-opaque spheres (Boone *et al.*, 1991; Eldridge *et al.*, 1996; Schreiner *et al.*, 1997) are that it allows arbitrary C-arm poses and that the fiducials do not completely obscure any part of the image, although the image contrast is reduced.

For extrinsic calibration, our strategy relies on identifying points and lines within a single fluoroscopic X-ray image of a known spatial arrangement of fiducials and computing the camera transformation parameters. This is simpler and more effective than using a calibration object holding radio-opaque spheres arranged in a spiral pattern (Eldridge *et al.*, 1996) or including fiducial lines in the bone fixator. Having the robot hold a calibration object either inside the femoral cavity or very close to the bone provides a common reference to co-register the intraoperative fluoroscopic images and to obtain directly the transformation between the robot's end-effector and the patient's femur, expressed in robot coordinates. We have integrated radio-opaque features into a radiolucent calibration object that can either be built into the surgical cutter bearing sleeve or else can be quickly mounted directly on the surgical cutter. The geometries are designed to permit the estimation of the relative position and orientation from single 2-D fluoroscopic X-ray images. Four fluoroscopic images—two for the dewarp grid and two roughly orthogonal for calibration—are the minimum number of images that were deemed to be sufficient to obtain the desired accuracy results.

Figure 8 shows the set-up of a positioning experiment performed with the Johns Hopkins University/IBM LARS robot (Eldridge *et al.*, 1996) and a C-arm. The LARS robot held a ball probe and was positioned in the interior of a radiolucent box into which 1/8 in steel balls had been inserted at known positions. The inside surface of the box was covered with copper foil and electrical contact between the foil and the ball probe could be detected by the LARS controller. After initial calibration to determine the transformation between the robot's tool holder and the calibration object, the robot was guided manually so that the ball was inside the box. Then, fluoroscopic X-ray images were taken from two roughly orthogonal orientations. After locating the coil and balls in

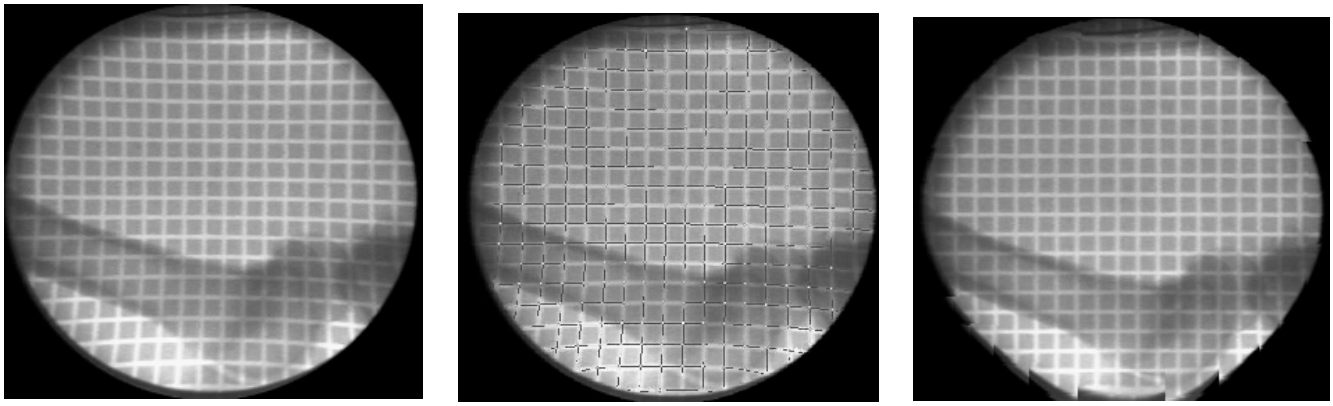

(a) (b) (c)

Figure 7. Illustration of the dewarping process: a fluoroscopic X-ray image of a smoked ham and a dewarping calibration plate: (**a**) before dewarping; (**b**) after groove location; and (**c**) dewarped. The white lines correspond to grooves cut into an aluminum dewarping calibration plate mounted on the X-ray detector.

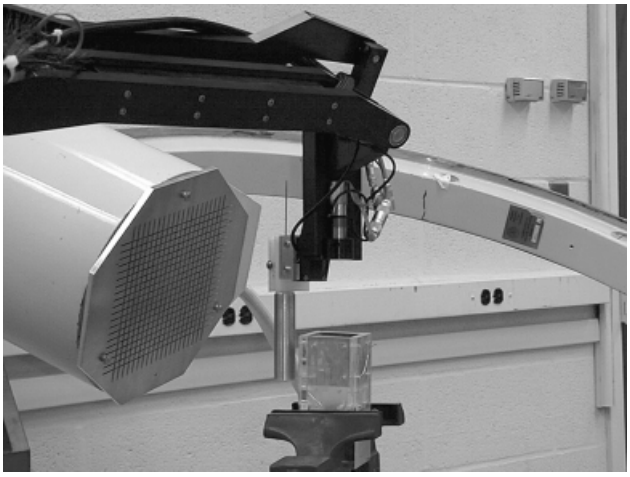

Figure 8. Experimental set-up at The Johns Hopkins University for image-guided positioning showing the C-arm with the dewarping grid plate mounted on its image intensifier (left), the LARS robot arm holding the probe (center top) and a box (center bottom).

both images, the relative difference in position and orientation of the C-arm between the two views was computed, and the spatial positions of the balls relative to the coil were estimated by triangulation. The positions were then used to compute the pose of the box relative to the robot. Finally, the robot located the interior surfaces of the box by moving in desired directions until electrical contact was made. Comparing the observed distances moved to the distances predicted from the images, we measured errors that were <1 mm for motions of up to 25 mm, and <0.3 mm for motions of 6 mm or less. These encouraging preliminary results suggest that the progressive refinement strategy, consisting of initial undercutting followed by re-registration before the finish cuts are made, is indeed viable.

5.4. Fluoroscopic X-ray and CT image contour extraction

We have developed a method for extracting contours from fluoroscopic X-ray and individual CT slice images, and a method for extracting surfaces from a set of CT slices. While both problems have been studied extensively in the literature, no robust, reliable and fully automatic method is yet available. Extracting bone contours reliably from fluoroscopic X-ray images is difficult because the images are noisy, have limited resolution, exhibit non-uniform exposure variation across the field of view and have varying contrast and exposure from shot to shot. The bone structures are surrounded by tissue, contain overlapping contours and have internal contours. In the CT data, the individual 2-D image quality is significantly better, but the spacing between slices can be significant (up to 6 mm), making it difficult to use iso-surface voxel techniques to extract the surfaces directly.

To find bone surface contours in fluoroscopic X-ray images and individual CT slices, we use a modified, semi-automatic active contour models technique (Kass *et al.*, 1987). In this method, an initial user-defined spline contour is stretched and bent according to an energy function acting through a potential field. The potential field is defined by the norm of the image intensity gradient. We use a smoothed image gradient norm for the potential and an initial polygonal contour model defined by a few user-supplied points in the vicinity of the structure of interest. The initial contour is deformed and fitted to the image data by minimizing its

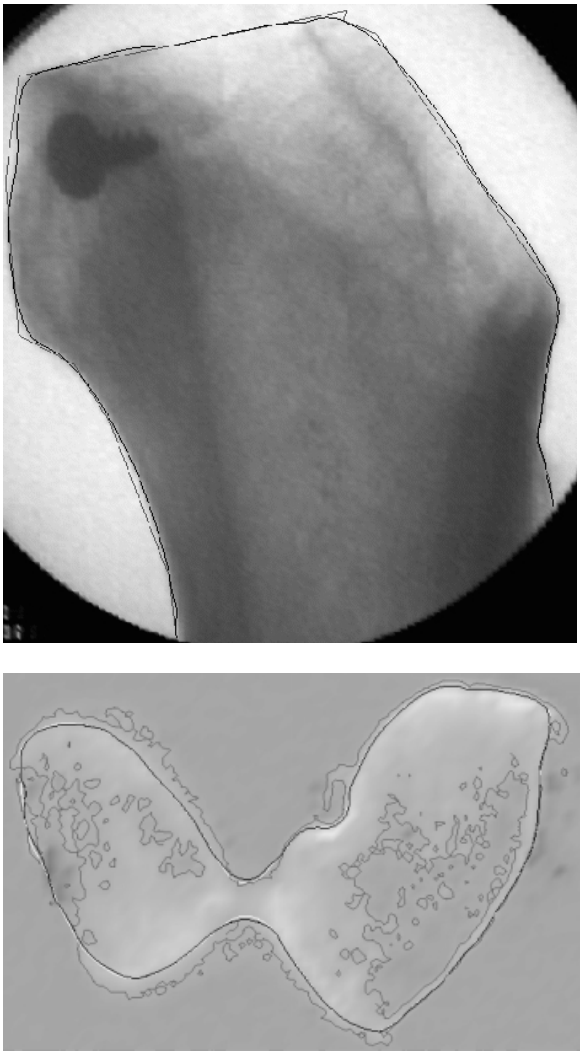

Figure 9. Extracted contours of a fluoroscopic X-ray image (top, black pixels indicate the extracted bone contours), and extracted contours from distal femur CT image slice (bottom). The continuous line (black pixels) is the contour obtained by our active contour method. The spurious gray contours is the best attempt produced by iso-contour-based segmentation and show that this method is inadequate even for a single CT slice.

energy by solving a partial differential equation using the finite-difference method. To define a longer range potential, we have implemented a two-step process. First, the initial deformable model is attracted by a smoothed low-resolution potential. Then, the result is used as the initial model for a second deformation step at full potential resolution. Because the contour topology is relatively simple, consisting of at most two disconnected contours, this method yields good

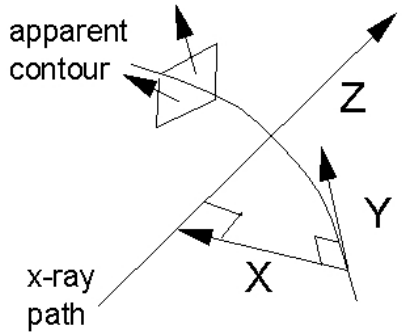

Figure 10. A directed X-ray path 'pierces' an apparent contour if the three vectors X, Y and Z form a right-handed frame. Vector X is the shortest distance between the bone contour and the X-ray path, Y is the oriented tangent of the bone contour and Z is the direction of the X-ray path.

results. For CT slices, we use the points in the neighboring slice as initial estimates for the following slice. Figure 9 shows examples of extracted contours in a fluoroscopic X-ray image and in a CT slice.

To construct a surface model from a set of CT slice images, we connect the extracted contours in successive CT slices by creating a triangular mesh. The mesh is constructed by finding corresponding points on the upper and lower contours and connecting them with an edge. Finding the correspondences is relatively straightforward because the femur anatomy is simple and known. The algorithm builds a surface patch connecting the two contours using a variation of the method described in Fuchs *et al.* (1977). The algorithm maximizes the sum of the quality measures (e.g. the ratio of the positive area by the sum of squared side lengths) of each surface triangle. The resulting surfaces are simplified with a tolerance between 0.0 and 0.3 mm using the variable tolerance method (Guéziec, 1995). Figure 11 show details of the constructed surface of the proximal and distal portion of a femur.

5.5. Anatomy-based matching of fluoroscopic X-ray images to CT datasets

We have developed a new method for matching fluoroscopic X-ray images to CT datasets. The method finds a rigid transformation optimizing the position and orientation of the bone surface such that its projections on the fluoroscopic X-ray images at the computed camera locations best match the bony anatomy observed in such images.

The fluoroscopic C-arm camera can be modeled as a pinhole projective camera with known focal length. The

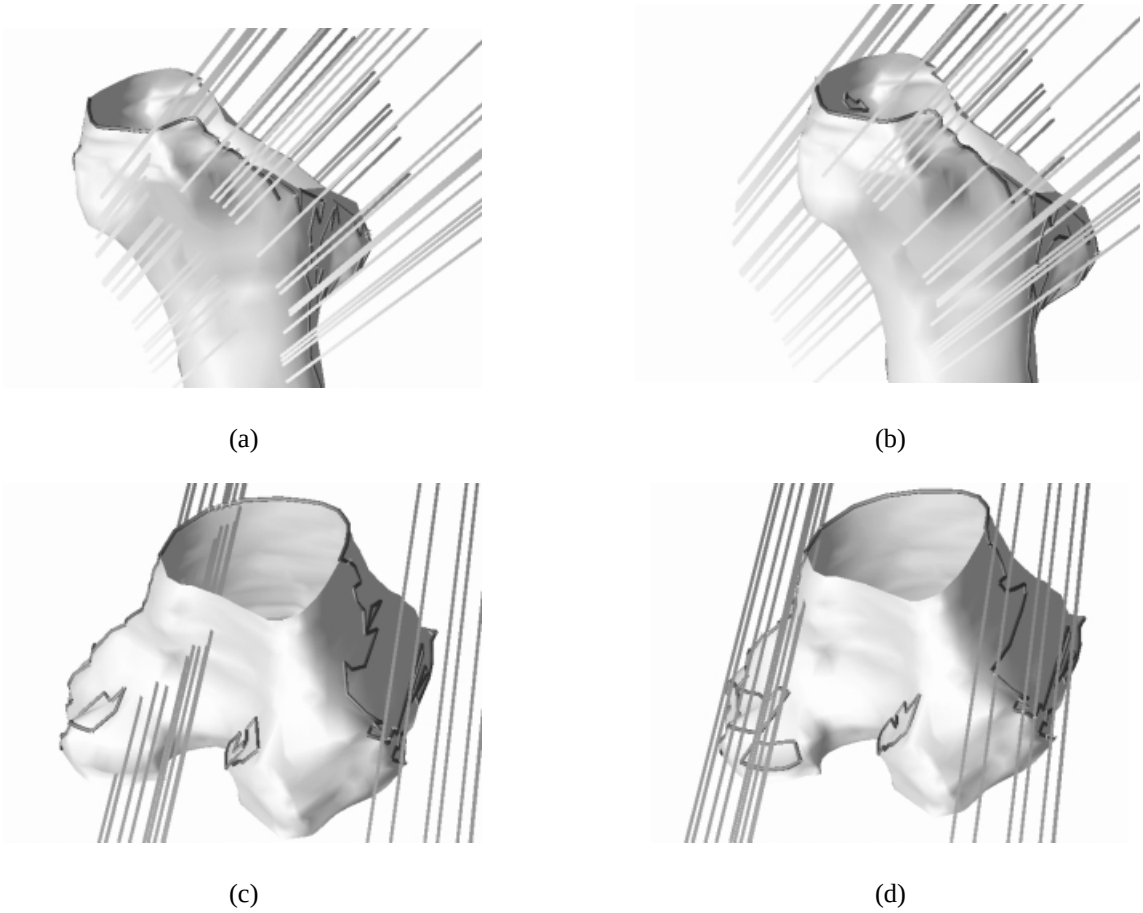

Figure 11. Proximal and distal femur surface models and X-ray paths: (**a**), (**c**) before registration and (**b**), (**d**) after registration using one proximal marker in addition to the fluoroscopic X-ray images.

pinhole center defines a set of lines emanating from it: when the bone surface model pose corresponds to the imaging pose, the lines emanating from the pinhole and grazing the bone surface correspond to points on the bone contour, as seen on the fluoroscopic X-ray images. We call the contour defined by the lines the apparent contour (Figure 10). A bone surface model which is at a location that does not match the bone contour in the X-ray camera pose can be brought closer to it by minimizing the distance between the apparent contour and the line bundle (Figure 11).

The iterative registration proceeds in four steps:

(i) compute the apparent contour on the bone surface;

(ii) for each line, find the closest apparent contour point;

(iii) compute, after eliminating outliers, a rigid transformation (rotation and translation) that minimizes the sum of squared distances between the apparent contour points and the lines; and

(iv) apply the transformation to the bone surface.

This process is repeated until the magnitude of the incremental rotation and translation is below a pre-specified threshold or the maximum number of iterations is reached. We describe each step briefly next (see Guéziec *et al.*, 1998 for details).

The bone surface apparent contours for a given C-arm camera pose are computed as follows. For each surface triangle, the viewing direction is defined as the vector originating from the pinhole center to the triangle centroid. The triangle is said to be visible if its normal, defined by the cross product of ordered oriented triangle edges, makes an obtuse angle with the viewing direction, otherwise it is invisible. Surface apparent contours are the subset of surface edges such that the triangle on one side of the edge is visible and the triangle on the other side of the edge is invisible. The apparent contours are constructed by linking these edges to form non-planar closed polygonal curves in space. The algorithm builds the apparent contour by first identifying the

edges belonging to it and orienting them so that the visible triangle is to the left of the edge. It then picks an edge and adds edges of neighboring triangles iteratively until the circuit is closed. This process is repeated on the remaining edges until all apparent contours are found.

The closest points of the apparent contours to each line are computed by first computing the closest point from the line in all apparent contours and then selecting a particular apparent contour and its corresponding closest point. For the first step, we use a hierarchical decomposition of the polygonal apparent contours into segments with an associated bounding region that completely encloses the contour portion of the segment. The hierarchical decomposition allows us to find a closest point in $O(\log n)$ average expected time, where n is the number of edges in the apparent contour. For the second step, the algorithm tests whether the line 'pierces' the apparent contour (Figure 10). If so, the algorithm keeps the pierced apparent contour for which the distance to the closest point is the largest. Otherwise, it keeps the one for which the distance to the closest point is the smallest.

To find the transformation matrix that minimizes the sum of squared distances between the apparent contour points and the lines, we use the following formulation:

$$\min_{\mathbf{Q},\mathbf{t}} \sum_{i=1}^{m} d_i^2 = \min_{\mathbf{Q},\mathbf{t}} \sum_{i=1}^{m} \|\mathbf{V}_i(\mathbf{c}_i - (\mathbf{Q}\mathbf{p}_i + \mathbf{t}))\|^2 \qquad (1)$$

$$\mathbf{V}_i = \begin{bmatrix} 0 & -v_3 & v_2 \\ v_3 & 0 & -v_1 \\ -v_2 & v_1 & 0 \end{bmatrix} \qquad (2)$$

where $\mathbf{p}_i$ are closest apparent surface points, $\mathbf{c}_i$ is the camera pinhole position, $\mathbf{v}_i$ are the unit direction vectors of the lines, and $\mathbf{Q}$ and $\mathbf{t}$ are the unknown rotation matrix and translation vector, respectively. We use the Cayley rotation parameterization, which states that if $\mathbf{U}$ is a skew-symmetric matrix obtained from a vector $\mathbf{u}$, then the matrix

$$\mathbf{Q} = (\mathbf{I} - \mathbf{U})(\mathbf{I} + \mathbf{U})^{-1} \qquad (3)$$

is a rotational matrix. The advantage of this form is that for small rotations, $\mathbf{Q} \sim \mathbf{I} - \mathbf{2U}$. The problem becomes

$$\min_{\mathbf{u},\mathbf{t}} \sum_{i} \|\mathbf{V}_i(\mathbf{p}_i - \mathbf{c}_i + 2\mathbf{P}_i\mathbf{u} + \mathbf{t})\|^2 \qquad (4)$$

where $\mathbf{P}_i$ is the skew-symmetric matrix derived from $\mathbf{p}_i$. Following Kumar (1992), we use the Tukey weighting function to scale each distance error. After some additional computations detailed in Guéziec *et al.* (1998), we obtain a linear system that we solve for the desired transformations.

We have conducted a series of experiments to assess the accuracy of our method. In one experiment, we used implanted fiducials on a cadaver bone and ROBODOC to obtain precise and reliable transformations that can be compared with the image-based method. We found that the femur is appropriate for 2-D/3-D registration because the images show more than just the femur shaft. The condyles and the anatomical features that are left in the proximal part of the femur provide sufficiently unique asymmetric and unique features for matching. A measured registration error of 2.3 mm was recorded in the dataset of Figure 11. The error is the maximum of the distances between the geometric location of the pins as measured by the robot and the computed ones.

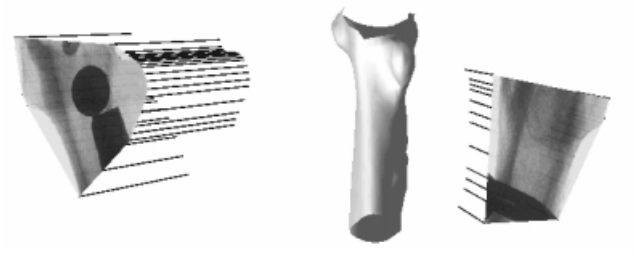

Figure 12. Experimental result of pinless X-ray image-based registration. The image shows a proximal femur surface model built from CT data (center) and two distortion-corrected X-ray fluoroscopic images (left and right) after registration in their original positions. The dark line segments emanating from the X-ray fluoroscopic image bone contours represent X-ray paths.

In a second experiment (Figure 12), we used a different calibration rod and no pins and obtained a maximum registration error measured at each marker location of between 1.2 and 3.6 mm. The main cause of the error is the effect of angular errors on a long shaft: an angular error of 1°, which is common in such applications, on a femur 450 mm long causes a displacement of 8 mm at the tip or 4 mm if the center of rotation is in the middle of the femur.

5.6. Image spreadsheet

We are developing an image spreadsheet for selectively viewing X-ray images, CT cross-sections and 3-D volumetric reconstructions (Figure 13). The spreadsheet maintains the images co-registered and allows manipulating overlays on them. It includes standard image-processing tools, such as histogramming, intensity adjustments, and zooming and panning. It allows the user to specify the number of windows desired (four in Figure 13) and maintains a scrollable window (bottom window) containing thumbnail views of the case images, which can be displayed by dragging and dropping them in any window. New images can be generated, saved and added as thumbnail images to the bottom window. Volume data, such as CAD implant models and cut volumes can be overlayed on bitmap images.

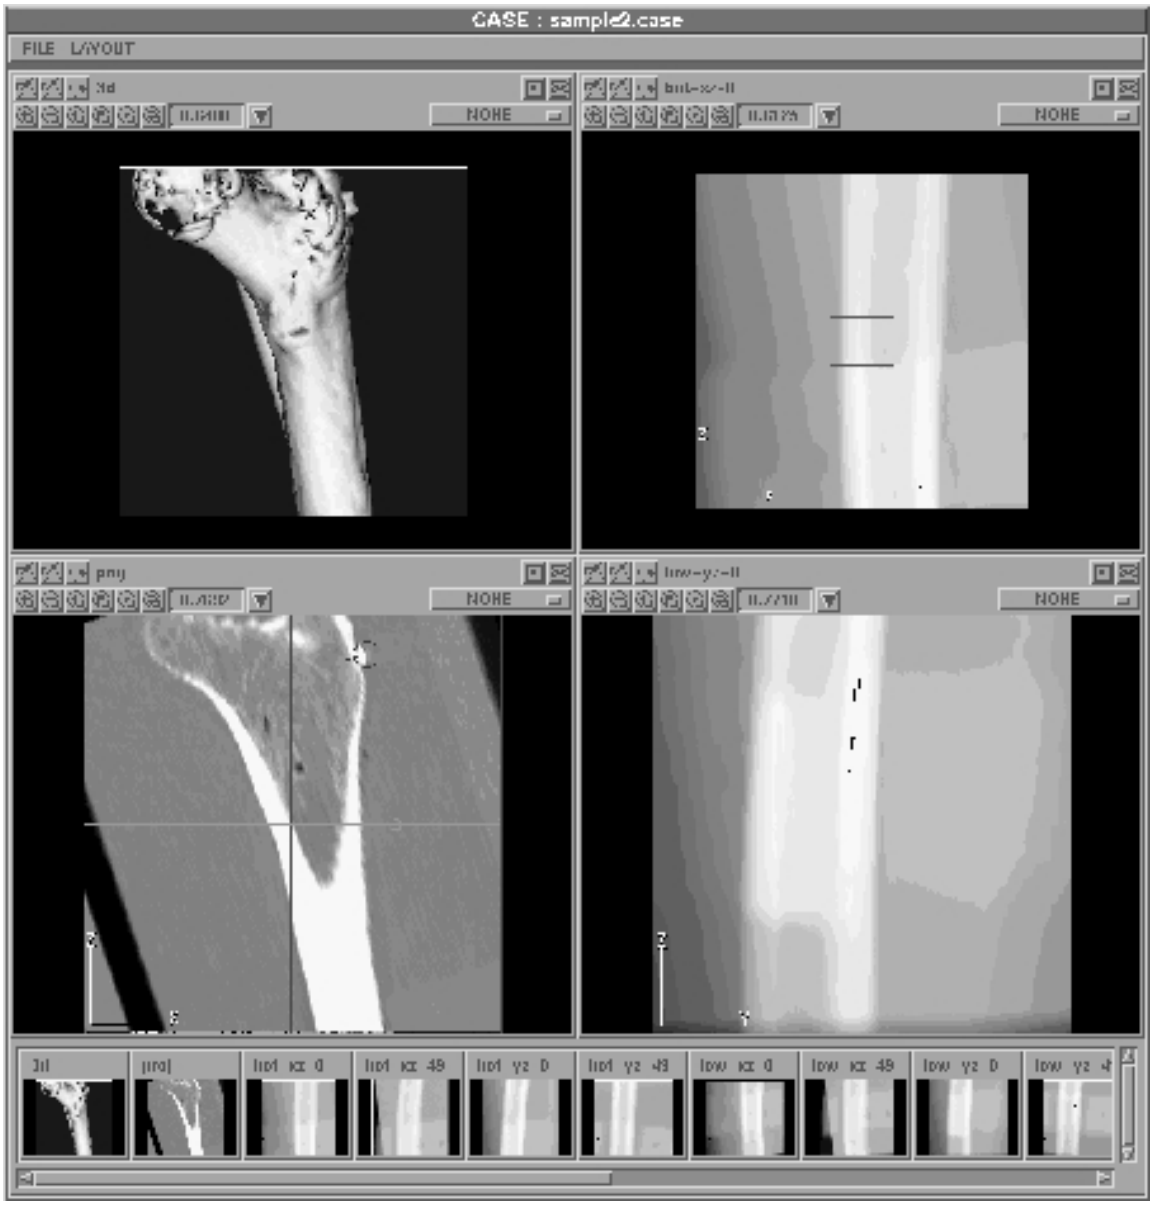

Figure 13. Image spreadsheet with four windows and thumbnails.

The user can directly manipulate and position the implant and the volumes with the mouse. The spreadsheet algorithm computes the projection and keeps them co-registered. Another alternative, not yet implemented, is semi-automatic positioning using surgeon-defined correspondences. The surgeon can use the mouse to designate points on the implant that should align with points in individual X-ray and CT images. The system then computes an implant position and orientation that brings the selected pairs of points as close together as possible by formulating and solving a least-squares minimization problem, achieving an optimal placement with respect to the specified correspondences. By interactively adding, deleting and modifying correspondences, the surgeon can quickly find the best implant or volume position. This method is potentially less time consuming because it reduces divergences on several individual views

simultaneously. However, it requires the surgeon to get used to a different kind of positioning, which has not yet been found necessary by ORTHODOC users.

5.7. Cement machining

We have conducted several experiments to assess the cement removal process. In one experiment, we tested whether the cutters currently used in ROBODOC PTHR surgery are adequate to cut bone cement by cutting circular shapes in a hard plastic material with a density similar to that of bone cement. To determine accuracy, the diameter of the cavities was measured and compared with the planned diameter, obtaining satisfactory results for shape and position accuracy. In another experiment, we tested how deep we can cut in bone cement and still achieve the accuracy needed. With the current instrumentation, the ROBODOC system can cut an implant cavity $\sim$300 mm deep along the axis of the bone. We found that proper irrigation with a sterile saline solution is very important for cement machining. Full quantitative evaluation is part of future work.

5.8. Clinical RTHR surgery experience with ROBODOC

Two surgeons at the Berufsgenossenschaftliche Unfallklinik (BGU) in Frankfurt, Germany, conducted four ROBODOC RTHR surgeries at the end of 1996 under the supervision of ISS engineers. The system used was the PTHR ROBODOC system augmented with imaging metal artifact reduction and interactive cut-volume definition software (Figure 5). The cases were carefully chosen to be within the specifications of the extended system.

The surgeries were successful in all four cases, and demonstrated the need for improvements in the cement cutting procedure. With the original ROBODOC PTHR surgery settings, the cutter had great difficulty machining the cement, causing frequent 'force freezes', i.e. excessive force when cutting, which causes the cutting to stop during the procedure. In the first surgery, there were 73 force freezes, requiring a cement cutting time of over 2 h (versus the expected 0.5 h). To improve the cement cutting performance of the robot, the ISS team modified the cavity generation software to reduce the amount of cement that the robot had to machine at each pass and improved the recovery capability from the force freezes. The second surgery required slightly over an hour of cement cutting time, with only 29 force freezes. During the third surgery, the robot machined the proximal portion of the cement mantle for $\sim$10 min, at which time the surgeons noticed that the whole mantle had come loose and removed it manually. This was considered a very good clinical result: it is likely that the robot loosened the cement mantle by machining away the proximal part. By the fourth case performed at the beginning of 1997, the software user interface had improved to the point where ISS engineers were involved only in the planning phase remotely via modem.

To allow the surgeons to perform all parts of the planning without ISS assistance, the following software improvements were made.

- Added a graphical indication of the minimum volume required for cutter clearance on each cross-section. Without this indication, the system will fail to create a cavity without a meaningful error message if the user created a contour that was too small for the cutter.
- Allow the user to insert new contours in between existing contours. This was necessary if the surgeon spaced the initial contours too far apart to create an accurate cavity.
- Improve the reliability of the cavity generation software so that it works in $>$90% of the cases, and will fail gracefully otherwise, i.e. create the correct cavity or no cavity at all.
- Create a tool for the surgeon to mark the top of the bone. This information is passed to the cavity generation software so that it can generate proper cutter paths for the 'precut' area, i.e. the bone proximal to the cement cavity.

The improved software was tested successfully on April 17, 1997 at BGU where the surgeons performed the first revision case without any involvement from ISS. Since then, the software has been used by surgeons at BGU for several RTHR procedures. The primary limitation of the current system is that the femoral fixator cannot be used for weak bones (as is often the case in revision patients) because it can damage the bone. We are currently investigating new fixator designs that can accommodate weaker bones.

6. CONCLUSION

We have presented the concept and preliminary results for a computer-integrated system to assist surgeons in revision total hip replacement (RTHR) surgery. Our starting points were the ROBODOC system for primary hip replacement surgery and the manual RTHR surgical protocol. We identified four key issues that must be addressed in developing the system:

(i) extended preoperative planning to reduce imaging metal artifacts in CT data;

(ii) interactive cement cut-volume definition;

(iii) intraoperative plan validation and modification; and

(iv) image, patient and robot registration.

We described methods for dealing with the above problems, and results of experiments for better understanding the technical feasibility of alternative solutions. Clinical trials were conducted with an extended RTHR version of ROBODOC, which included CT imaging metal artifact removal and interactive cement volume definition. Our preliminary results indicate that the proposed approach is viable and could form the basis of a practical system for clinical use. We believe that some of the solutions developed specifically for RTHR surgery will be applicable to many orthopedic and other surgical problems.

The imaging metal artifact reduction algorithm has proved to be of great use and will be incorporated into ISS' ORTHODOC product. Version 1.0 of the RTHR system has just been released. Short-term plans include the design of a new femoral fixator that can handle weaker bones, the combination of the cement and implant cavity machining (currently, the system only machines cement), and automating the method for identifying the cement cut volume. Further improvements and clinical validation of fluoroscopic registration methods are required.

ACKNOWLEDGEMENTS

Janez Funda, Steve Gomory, Gabriel Taubin, Roger Kane, Lisa Brown, Joe McCarthy MD, Roderick Turner MD and Bela Musits participated in the early stages of this project. This project is partially funded by the US Department of Commerce under the National Institute of Standards' Advanced Technology Program (NIST/ATP) Cooperative Agreement Number 94-01-228. The Johns Hopkins University members of our team also gratefully acknowledge generous donation of computing equipment and software to Johns Hopkins by Intel Corporation and Microsoft Corporation. We thank the anonymous reviewers for suggestions that helped improve the presentation of the paper.

REFERENCES

(1992) *Orthopedic Network News* Vol. 3–4, No 4. Mendenhall Associates Inc.

Bookstein, F. L. (1991) *Morphometric Tools for Landmark Data*. Cambridge University Press, Cambridge.

Boone, J. M. *et al.* (1991) Analysis and correction of imperfections in the image intensified-TV-digitizer imaging chain. *Med. Phys.*, 18, 236–242.

Börner, M., Lahmer, A. and Stier, U. (1998) Experiences with the ROBODOC system in more than 1000 cases. In Lemke, H. *et al.* (eds), *Proc. 12th Int. Symp. on Computer-Aided Radiology and Surgery*, pp. 689–697. Elsevier, Amsterdam.

Brack, C., Burghart, R., Czpof, A., *et al.* (1998) Accurate X-ray navigation in computer-assisted surgery. In Lemke, H. *et al.* (eds), *Proc. 12th Int. Symp. on Computer Assisted Radiology and Surgery, Berlin*. Springer-Verlag, Berlin.

Brandt, G., Rademacher, K., Lavallée, S., Staudte, H.-W. and Rau, G. (1997) A compact robot for image-guided orthopaedic surgery: concept and preliminary results. In Troccaz, J. *et al.* (eds), *Proc. 1st Joint Conf. on Computer Vision, Virtual Reality and Robotics in Medicine and Medical Robotics and Computer-Assisted Surgery*, Grenoble, *Lecture Notes in Computer Science*, Vol. 1205, pp. 767–776. Springer-Verlag, Berlin.

Brown, L. M. (1996) Registration of planar film radiographs with computed tomography. *Proc. IEEE Conf. on Mathematical Methods in Biomedical Image Analysis*.

Browne, J. A., Herman, G. T. and Odhner, D. (1993) SNARK93: a programming system for image reconstruction from projections. *Technical Report* MIPG198, Department of Radiology, University of Pennsylvania, 12, pp. 236–241.

Crenshaw, A. H. (ed.) (1987) *Campbell's Operative Orthopaedics*. Mosby Year Book.

Champleboux, G., Lavallée, S., Sautot, P. and Cinquin, P. (1992) Accurate calibration of cameras and range imaging sensors, the NPBS method. *Proc. IEEE Int. Conf. on Robotics and Automation*, Nice. IEEE Press, Piscataway, NJ.

Duchon, J. (1976) Interpolation des fonctions de deux variables suivant le principe de la flexion des plaques minces. *Revue Francaise d'Automique, Informatique, Recherche Operationnelle (RAIRO) Analyse Numerique*, 10, 5012.

Eldridge, B., Gruben, K., LaRose, D., Funda, J., Gomory, S., Karidis, J., McVicker, G., Taylor, R. H. and Anderson, J. (1996) A remote center of motion robotic arm for computer assisted surgery. *Robotica*, 14, 103–109.

Fuchs, J., Kedem, Z. M. and Uselton, S. P. (1977) Optimal surface reconstruction from planar contours. *Commun. ACM*, 20, 693–702.

Glover, G. H. and Pelc, N. J. (1981) An algorithm for the reduction of metal clip artifacts in CT reconstructions. *Med. Phys.*, 8, 799–807.

Guéziec, A. (1995) Surface simplification with variable tolerance. *Proc. 2nd Int. Symp. on Medical Robotics and Computer Assisted Surgery*, Baltimore, MD, pp. 132–139. Wiley, New York.

Guéziec, A., Kazanzides, P., Williamson, B. and Taylor, R. H. (1998) Anatomy based registration of CT-scan and X-ray fluoroscopy data for intra-operative guidance of a surgical robot. *IEEE Trans. Med. Imag.*, 17, 715–728.

Hamadeh, A., Lavallée, S. and Cinquin, P. (1998) Automated 3-dimensional computed tomographic and fluoroscopic image registration. *Comp. Aided Surg. J.*, 3, 11–19.

Hinderling, T., Ruegsegger, P., Anliker, M. and Dietschi, C. (1979) Computed tomography reconstruction from hollow projections: an application to *in vivo* evaluation of artificial hip joints. *J. Comput. Assis. Tomogr.*, 3, 52–57.

Hofstetter, R., Slomczykowski, M., Bourquin, I. and Nolte, L. P. (1997) Fluoroscopy based surgical nagivation—concept and clinical applications. In Lemke, H. U. *et al.* (eds), *Proc. 11th Int. Symp. on Computer Assisted Radiology and Surgery*, Berlin. Springer-Verlag, Berlin.

Joskowicz, L. *et al.* (1995) Computer-integrated revision total hip replacement surgery: preliminary report. *2nd Int. Symp. on Medical Robotics and Computer Assisted Surgery*, Baltimore, MD, pp. 193–202. Wiley Publishers, New York.

Kalvin, A. and Williamson, B. (1998) Using scout images to reduce the effects of metal artifacts in CT. In *Medical Imaging 1997: Image Processing*, Newport Beach, CA, February. *Proc. Soc. Photo-Optical Instrum. Engng*, 3034, 1017–1028.

Kass, M., Witkin, A. and Terzopolous, D. (1987) Snakes: active contour models. *Int. Conf. on Computer Vision, Graphics, and Image Processing*, London, pp. 259–268. IEEE, Piscataway, NJ.

Kalender, W. A., Hebel, R. and Ebersberger, J. (1987) Reduction of CT artifacts caused by metallic implants. *Radiology*, 164, 576–577.

Klotz, E., Kalender, W. A., Sokiranski, R. and Felsenberg, D. (1990) Algorithms for the reduction of CT artifacts caused by metallic implants. *Medical Imaging IV: PACS System Design and Evaluation*, Newport Beach, CA, Vol. 1234, pp. 642–650. SPIE, Bellingham, WA.

Kumar, R. (1992) *Model Dependent Inference of 3D Information from a Sequence of 2D Images*. Ph.D. Thesis, University of Massachusetts, Amherst, Report No COINS-TR92-04.

Lavallée, S. (1996) Registration for computer integrated surgery: methodology, state of the art. In Taylor, R. H., Lavallee, S., Burdea, G. and Mösges, R. (eds), *Computer Integrated Surgery* pp. 77–98. MIT Press, Cambridge, MA.

Lavallée, S., Sautot, P., Troccaz, J., Cinquin, P. and Merloz, P. (1994) Computer assisted spine surgery: a technique for transpedicular screw fixation using CT data and a 3D optical localizer. *Proc. 1st Int. Symp. on Medical Robotics and Computer Assisted Surgery*, Pittsburgh, PA.

Medoff, B. P. (1987) Image reconstructions from limited data: theory and applications in computerized tomography. In Stark, H. (ed), *Image Recovery: Theory and Application*. Academic Press, New York.

Oppenheim, B. E. (1977) Reconstruction tomography from incomplete projections. In Ter-Pogossian, M. (ed), *Reconstruction Tomography in Diagnostic and Nuclear Medicine*, pp. 155–183. University Park Press, Baltimore, MD.

Paul, H., Mittelstadt, B. D., Bargar, W. L., Musits, B. L., Taylor, R. H., Kazanzides, P., Zuhars, J. F. and Williamson, B. (1992) A surgical robot for total hip replacement surgery. *Proc. IEEE Int. Conf. on Robotics and Automation*. IEEE, Piscataway, NJ.

Robertson, D. D, Weiss, P. J., Fishman, E. K., Magid, D. and Walker, P. S. (1988) Evaluation of CT techniques for reducing artifacts in the presence of metallic orthopedic implants. *J. Comput. Assis. Tomogr.*, 12, 236–241.

Schurman, D. J. and Maloney, W. J. (1992) Segmental cement extraction at revision total hip arthroplasty. *Clin. Orthopedics Relat. Res.*, 285, 158–163.

Schreiner, S., Anderson, J., Taylor, R. H., Funda J., Bzostek, A. and Barnes, A. (1997) A system for percutaneous delivery of treatment with a fluoroscopically-guided robot. *Proc. Joint Conf. of Computer Vision, Virtual Reality, and Robotics in Medicine and Medical Robotics and Computer Surgery*, Grenoble. Springer-Verlag, Berlin.

Taylor, R. H., Mittelstadt, B. D., Paul, H. A., Hanson, W., Kazanzides, P., Zuhars, J. F., Williamson, B., Musits, B. L., Glassman, E. and Bargar, W. L. (1994) An image-directed robotic system for precise orthopaedic surgery. *IEEE Trans. Robotics Automation*, 10, 261–275.

Turner, R. H. and Emerson, R. H. (1982) *Revision Total Hip Arthroplasty*. Grune and Stratton, Orlando, FL.

Tuy, H. K. (1992) Algorithm to reduce clip artifacts in CT images. *Medical Imaging VI: Image Processing Conf.*, Newport Beach, CA, pp. 635–643. SPIE, Bellingham, WA.

Viant, W. J., Phillips, R., Griffiths, J. G., Ozanian, T. O., Moshen, A. M. M. A., Cain, T. J., Karpinski, M. R. K. and Sherman, K. P. (1997) A computer assisted orthopaedic surgical system for distal locking of intramedullary nails. *Proc. Inst. Mech. Engng, Part H. (Journal of Engineering in Medicine)*, 211, No. H4, 293–300.

Wolberg, G. (1990) *Digital Image Warping*. IEEE Computer Society Press, Los Alamitos, CA.

Yaniv, Z., Joskowicz, L., Simkin, A., Garza-Jinich, M. and Milgrom, C. (1998) Fluoroscopic image processing for computer-aided surgery. In Wells, S. *et al.* (eds), *1st Int. Conf. on Medical Computing and Computer-Assisted Intervention, Lecture Notes in Computer Science*, Boston, MA, Vol. 1496, pp. 325–334. Elsevier, Amsterdam.