An Image-Directed Robotic System for Precise Orthopaedic Surgery

Russell H. Taylor, *Fellow, IEEE*, Brent D. Mittelstadt, Howard A. Paul, William Hanson, Peter Kazanzides, *Member, IEEE*, Joel F. Zuhars, *Member, IEEE*, Bill Williamson, Bela L. Musits, Edward Glassman, and William L. Bargar

Abstract—We have developed an image-directed robotic system to augment the performance of human surgeons in precise bone machining procedures in orthopaedic surgery, initially targeted at cementless total hip replacement surgery. The total system consists of an interactive CT-based presurgical planning component and a surgical system consisting of a robot, redundant motion monitoring, and man-machine interface components. In vitro experiments conducted with this system have demonstrated an order-of-magnitude improvement in implant fit and placement accuracy, compared to standard manual preparation techniques. The first generation system described in this paper was used in a successful veterinary clinical trial on 26 dogs needing hip replacement surgery. It was the basis for subsequent development of a second-generation system that is now in human clinical trials.

I. INTRODUCTION AND BACKGROUND

A. Augmentation of Human Skill in Surgery

THE RESEARCH reported in this paper represents a step in an evolving partnership between humans (surgeons) and machines (computers and robots) that seeks to exploit the capabilities of both to do a task *better* than either can do alone. Recent advances in medical imaging technology (CT, MRI, PET, etc.), coupled with advances in computer-based image processing and modelling capabilities have given physicians an unprecedented ability to model and visualize anatomical structures in live patients, and to use this information quantitatively in diagnosis and treatment planning. Further, advances in CAD-CAM technology have made it practical to use this data to design and precisely fabricate custom surgical implants for individual patients.

One result is that the precision of image-based presurgical planning often greatly exceeds the precision of surgical execution. Typically, geometrically precise surgery has been limited to procedures (such as brain biopsies) for which a suitable stereotactic frame is available. The inconvenience and restricted applicability of these devices has led many

researchers to explore the use of robotic devices to augment a surgeon's ability to perform geometrically precise tasks planned from computed tomography (CT) or other image data.

The pioneering work in the use of general purpose robots for surgery was that of Kwoh *et al.* [1] who used a six-axis industrial robot to replace a stereotactic frame in neurosurgery. In this case, the robot was mounted in a known position relative to the table of a CT scanner and suitable geometric calibrations were performed. During surgery, the patient was CT-scanned and a desired placement for a biopsy needle probe was determined from the image data. The robot then positioned a passive needle guide appropriately, brakes were applied, and power was turned off. Finally, the surgeon inserted the needle through the guide into the patient's brain. The principal benefit gained was the greater convenience and faster positioning possible with the robot, compared to the use of a stereotactic frame. A number of similar systems have been developed subsequently. The most successful to date is that of Lavallee, *et al.*, ([2], [3]), who used a stereo pair of intraoperative radiographs to register the robot to the patient's CT data (and to the patient) and to plan needle paths that avoid blood vessels. Over three hundred cases have been performed, although (again) the robot is turned off while the needle is inserted. Kelly, *et al.* [4], [5] have implemented a specialized motorized stereotactic system for laser neurosurgery, in which an XYZ table is used to reposition the patient's head relative to the focal point of a surgical microscope. More recently, Drake, Goldenberg, *et al.* [6] have reported several cases in which a general purpose robot moved while in contact with the patient, although the motions were very simple and highly constrained. Further, these cases were performed on an exception basis, in which the surgeon had no practical alternative, despite somewhat more limited safety checking than would have been desirable for more routine use. Several other neurosurgery robots are in various stages of development (e.g., [7]).

A number of active robotic systems for augmentation of non-neurosurgical procedures have also been proposed or developed. For example, Davies *et al.* have developed a specialized robotic device to assist in laparoscopic prostatectomies [8], which has been used clinically. A number of groups (e.g., [9]–[12]) have developed a variety of other telerobotic devices for endoscopic and laparoscopic surgery. McEwen *et al.* have developed and marketed a clinically qualified voice controlled limb positioning system for orthopaedics [13]. Several groups (e.g., [14], [15]) have demonstrated *in vitro* robotic systems for

positioning passive instrument guides for knee replacement surgery. Of these applications, that of Davies comes closest to ours in the sense that it uses an active automatic device to perform a tissue removal operation. Important differences include an order-of-magnitude difference in the accuracy required for the application, the greater complexity of the shapes to be cut, the use of a general purpose manipulator rather than a specialized devise, and the greater degree of safety and consistency checking built into our system, which must move safely in a much less constrained volume.

B. Precise Orthopaedic Surgery

Orthopaedic applications represent a particularly promising domain for the integration of image and model-based presurgical planning, CAD/CAM technology, and precise robotic execution. For example, about half of the 300 000 total hip replacement operations performed each year use cementless implants. In these procedures, accurate preparation of the femoral cavity to match the implant shape and accurate placement of the cavity relative to the femur can significantly affect stress transfer, implant stability, and restoration of proper biomechanics, which, in turn, are important factors affecting efficacy. For example, Sandbom, et al., [16]. have reported that the size of gaps between bone and implant significantly affects bone ingrowth. Furthermore, the present manual broaching method[1] for preparing the femoral cavity leaves considerable room for improvement. In one recent study Paul, Hayes, et al. [17] found that only about 20% of the implant actually touches bone when it is inserted into a manually broached hole. The average gap between the implant and the bone was commonly 1–4 mm and the overall hole size was 36% larger than the broach. Furthermore, the exact placement of the implant cavity relative to the bone (which affects restoration of biomechanics) depends on the surgeon's ability to line up the broach manually and to drive it the right distance into the femur. Driving the broach too far can split the femur.

These considerations have led us to explore the use of robotic machining to prepare the femoral cavity for the implant. Initial feasibility studies by Paul, Mittelstadt, et al. [18] demonstrated that a robot could successfully machine shapes in human cadaver bones and that preoperatively implanted calibration pins could be used to accurately register CT image and robot coordinates for a femur.

Following these studies, we developed a complete planning and execution system suitable for use in an actual operating room. In vitro experiments with this first generation system demonstrated an order-of-magnitude improvement in surgical precision, compared to manual broaching. One of the authors (Dr. Paul) conducted a veterinary clinical trial on dogs needing hip replacement surgery. This experience provided the basis for development of a second-generation system that is now in human clinical trials [19]–[22].

Subsequent sections of this paper will summarize the presurgical planning and surgical procedure followed for robotic

---

[1] Fig. 1 shows a typical cementless implant and the corresponding broach used to make the hole for it. Fig. 2 shows the use of a broach on a human patient. The procedure in a dog is essentially the same.
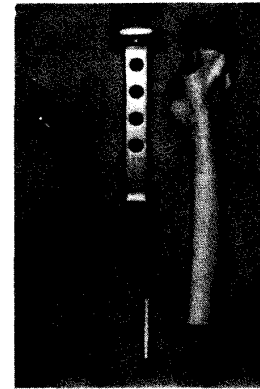

Fig. 1. Typical cementless hip implant and instrumentation. This figure shows a typical cementless hip implant, together with the broach used to produce a corresponding hole in the patient's thigh in conventional manual surgery. Proper placement of the implant socket relative to the femur and accurate reproduction of the socket shape are very important to assure stability, uniform stress transfer, and restoration of the proper biomechanics.

Fig. 2. Manual broaching procedure. This figure shows the use of a broach in a human cementless hip replacement. The procedure in a dog is essentially the same. One study found that only about 20% of the implant actually touches bone when it is inserted into a manually broached hole. The average gap between the implant and the bone was commonly 1–4 mm and the overall hole size was 36% larger than the broach.

hip replacement surgery and will discuss the requirements for robotic systems intended to augment human precision in surgery. After providing a brief overview of the system architecture, we will provide a fuller discussion of several key aspects of the system, including the image-based presurgical planning, geometric calibration, shape cutting and safety checking mechanisms. Finally, we will discuss experience of the system in actual clinical use (on dogs) and will discuss some of the lessons learned.

II. SUMMARY OF PROCEDURE

Before surgery, three titanium pins are implanted through small skin incisions into the greater trochanter and condyles of the patient's femur. A CT scan is made of the leg. The presurgical planning system automatically locates the pins relative to the coordinate system of the CT images. The sur-
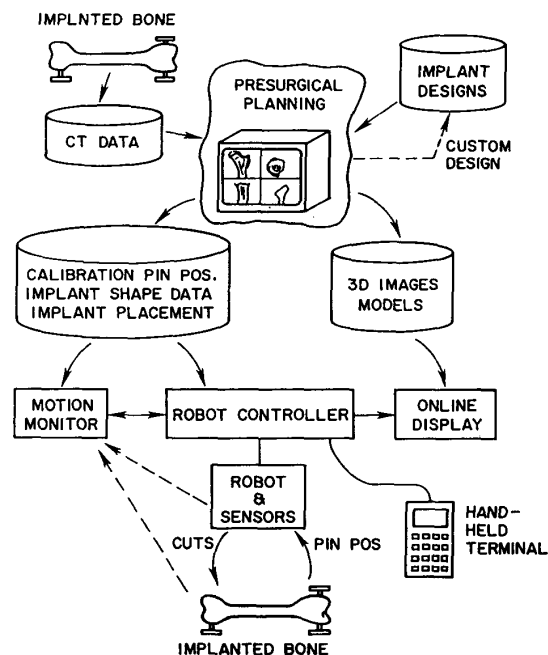
Fig. 3. Architecture of hip replacement surgery system. The system consists of a presurgical planning component and a surgical component. In the system used for the veterinary clinical trial, the motion monitoring and robot control functions are subsumed within the robot controller.

geon interactively selects an implant model and determines its desired placement relative to CT coordinates. This information is written to a diskette for use in surgery.

Key steps of the intraoperative procedure are shown in Fig. 5 for an *in vitro* test on a cadaver femur. Fig. 6 shows the operating room scene during the first canine clinical trial in May, 1990. Briefly, the procedure is as follows.

1) The robot is brought into the operating room and powered up. A sterile cutting tool is attached to a tool interface just below the force sensor, and the robot is covered with a sterile drape. The patient data diskette is loaded into the robot controller, and the robot is placed in a standby mode.

2) The patient is prepared and draped in the normal manner. Surgery proceeds normally until the acetabular component of the implant is implanted and the ball of the femur is removed.

3) The robot is brought up to the operating table, and the femur is rigidly attached to the robot base, using a specially designed fixator. The three titanium pins are exposed manually.

4) A ball probe "cutter bit" is inserted into the collet of the cutting tool. The top center of each pin is then located by a combination of manual guiding and autonomous tactile search by the robot. Although several modes of manual guiding are available, the most commonly used is force compliance. The surgeon simply pulls on the shaft of the cutter; the robot controller senses the forces exerted on the tool and moves the robot in the indicated direction.

5) The robot controller uses the pin location information to compute an appropriate transformation from CT coordinates to robot coordinates. The ball probe is replaced by a standard cutting bit, and the robot cuts out the desired implant shape at the planned position and orientation relative to the pins. The surgeon monitors progress both by direct observation of the robot and patient and by looking at a graphical display depicting successive cuts.

6) When cutting is complete, the femur is unclamped from the fixator, and the robot is moved out of the way. The rest of the procedure proceeds in the normal way, with the added step of removing the locator pins from the patient.

III. REQUIREMENTS AND ISSUES

A. Human-Machine Interaction in a Surgical Situation

Our goal is not to replace the surgeon. Instead, we are concerned with developing a surgical tool that can assist the surgeon by precisely executing a tissue removal task under the surgeon's supervision. Although the robot's geometric accuracy is much greater than the surgeon's, the surgeon's understanding of the total situation is clearly much greater than any computer's, and he or she is responsible for what goes on in the operating room. Suitable interfaces must be provided to allow the surgeon to monitor the robot's actions, to pause execution at any time, initiate error recovery actions, and provide positional guidance to the robot. There is also the related problem of human-computer interaction in presurgical planning. Convenient and naturally understood interfaces must be provided to allow the surgeon to specify what implant shape is to be cut and where it is to go. Furthermore, the interfaces used intraoperatively to report progress of the surgery should be as consistent as possible with those used to plan it.

B. Registration of Plan Data with Intraoperative Reality

The surgical plan is based on anatomical information derived from CT images taken prior to surgery. Reliable and accurate methods to locate the corresponding anatomical structures relative to the robot are essential if the plan is to be executed successfully.

C. Verification

It is very important to verify that the greater potential geometric accuracy offered by the use of a robotic surgical system is in fact achieved in practical use. Suitable methods must be developed for verifying the performance of individual system components and of the system as a whole.

D. Operating Room Compatibility and Sterility

It must be easy to incorporate the robot into a hospital's normal routine. It may be difficult for a hospital to dedicate an operating room to robotic surgery, and even if it does so it is important that maintenance not be disruptive.[2] Gener-

[2] These considerations led us to rule out some configurations (such as a Cartesian manipulator suspended from the ceiling) that might otherwise have been attractive.
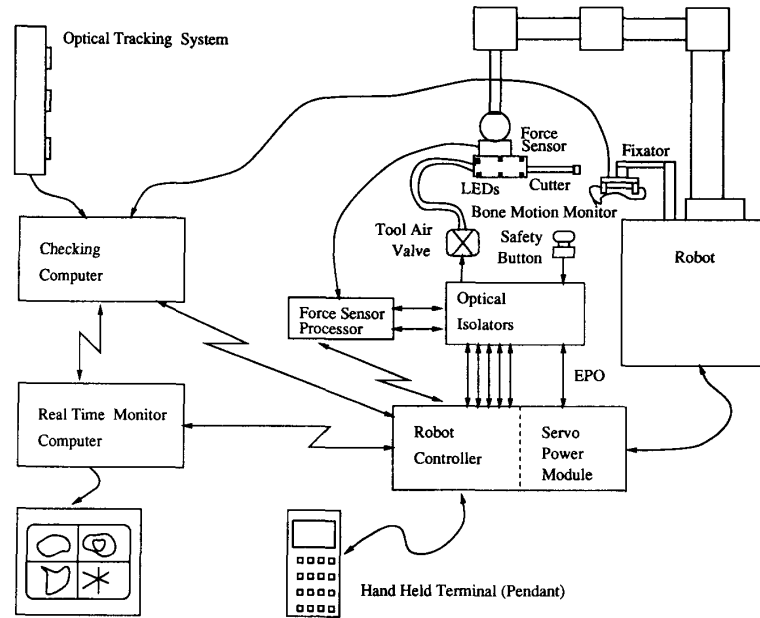
Fig. 4. Operating room system architecture. The operating room system consists of 1) a surgical robot with its associated controller, tooling, and safety interlocks, 2) a fixator to hold the bone securely to the robot, 3) a redundant motion-monitoring subsystem consisting of a checking computer, optical tracking system, and bone motion detector, and 4) a human-machine interface with an online display, display computer, and a hand-held terminal interfaced to the robot controller.

ally, the system should be easily brought into the operating room and set up as part of the normal presurgical routine. Similarly, removal, sterilization, and reattachment of the end effector and other critical components should be easy and suitable sterile drapes must be developed for the manipulator arm and other structural components that cannot be easily sterilized.

E. Safety, Error Recovery, and Backup

Clearly, redundant safety mechanisms are very important, both for the protection of the patient and of the surgeon. Manual pause and emergency power-off functions are essential. Wherever possible, potential error conditions must be anticipated and checked for, and adequate recovery procedures must be available. Although the robot often may be able to continue with the procedure following a pause, it is also prudent to provide a reliable means of stopping the robot, removing it from the surgical field, and continuing the operation with manual backup.

Since the surgeon must rely on the precision of the robot, it is extremely important that no single failure cause an undetected loss of accuracy. The system must monitor the position of the robot's cutting tool relative to the shape that it is supposed to cut and stop cutting if it strays out of the desired volume for any reason. It is especially important that systematic shifts (such as might arise from the bone slipping relative to the fixator) be detected promptly. A single misplaced cut can usually be repaired, but it may much harder to correct for misplacing the entire cavity.

IV. SYSTEM ARCHITECTURE

The system (Fig. 3) consists of a presurgical planning component and an intraoperative (surgical) component. These components are summarized below and discussed at greater length in subsequent sections.

A. Presurgical Planning

This component [23] implemented on an IBM workstation, permits the surgeon to select an implant model and size and to specify where the corresponding shape is to be machined in the patient's femur.

The system maintains a library of computer-aided-design (CAD) models of implant designs and accepts computed tomography data for individual patients. It automatically determines the CT coordinates of the preoperatively implanted locator pins and provides a variety of interactive graphics tools for the surgeon to examine the CT data, to select an appropriate model and size from the implant design library, and to manipulate the position and orientation of the selected implant shape relative to CT coordinates.[3] The output consists of files containing 1) patient identification data 2) the position of the locator pins relative to CT coordinates, 3) the implant specification, 4) the desired implant placement relative to CT coordinates, and 5) processed image and model data that will be used for a realtime animation of the progress of the surgery.

[3] In the future, we anticipate the use of computer optimization techniques to assist the surgeon in determining the best implant placement and also in the design of custom implants.
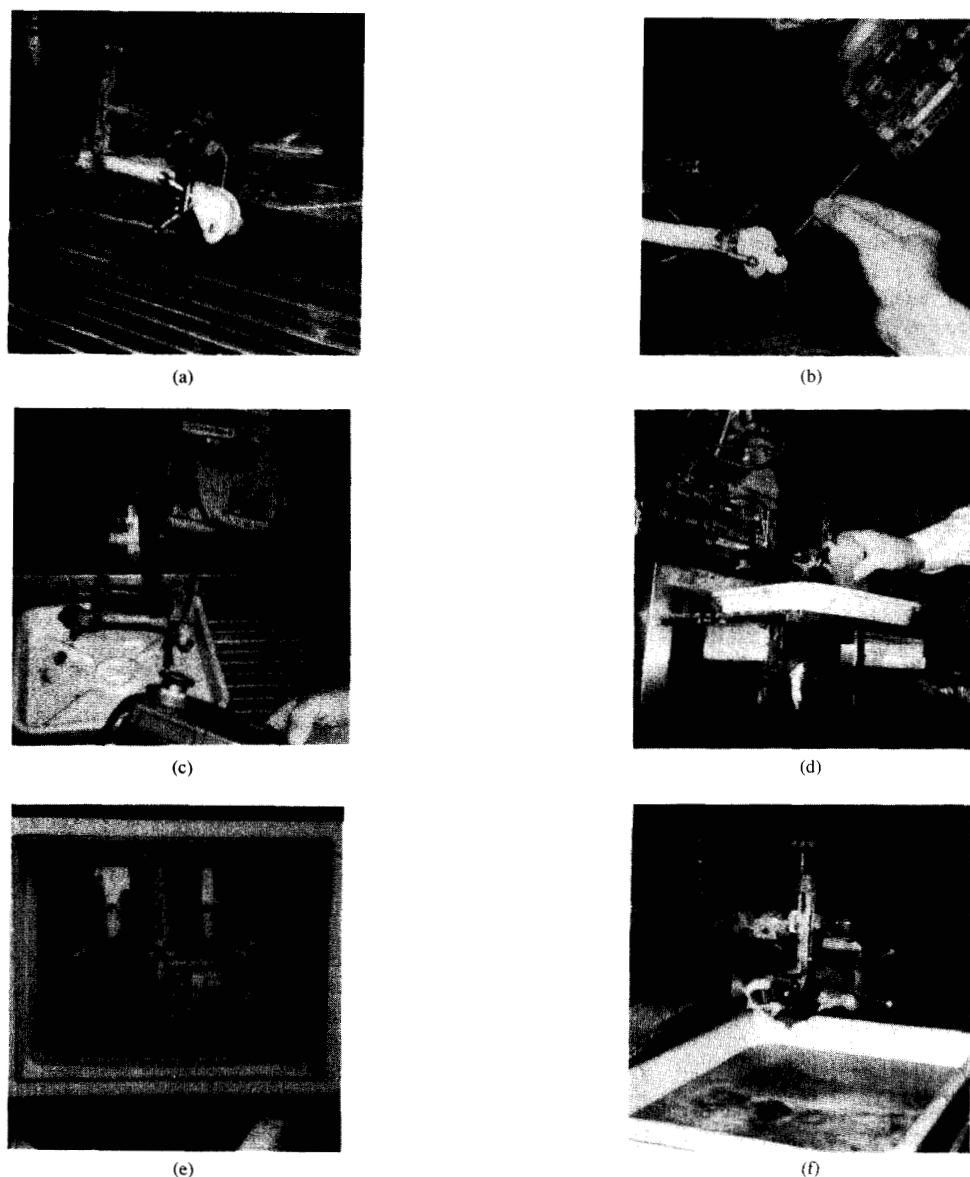
Fig. 5. Surgical procedure for hip surgery. (a) Fixated cadaver bone. (b) Manual guiding to approximate pin position. (c) Tactile search for a pin. (d) Cutting the shape. (e) Online display. (f) Final result.

B. Operating Room System

The operating room system (illustrated in Fig. 4) consists of several components. The five-axis *robot* is an IBM 7576 SCARA manipulator with an added pitch axis, six degree-of-freedom force sensor, and a standard high-speed surgical cutting tool. During surgery, all but the robot's cutting tool is covered by a sterile drape; the cutting tool is separately sterilized. A sterile *fixator* rigidly attached to the robot's base holds the bone during the robotic part of the procedure. The robot controller provides servocontrol, low-level monitoring, sensor interfaces, and higher-level application functions implemented in the AML/2 language. During surgery, the force sensor is used to support redundant safety checking, tactile search to find the locator pins, and compliant motion guiding by the surgeon.

The *redundant motion monitoring subsystem* [24] is implemented on an IBM PC/AT with specialized IO hardware. It relies on independent sensing to track the position and orientation of the robot end effector during the cutting phase of the surgery, and checks to verify that the cutter tip never strays more than a prespecified amount outside of the defined implant volume. It also monitors strain gauges that can detect possible shifts of the bone relative to the fixation device. If either condition is detected, a "freeze motion" signal is sent to the robot controller. After motion is stopped, application

code in the robot controller queries the motion monitoring system for more information and then enters an appropriate error recovery procedure under the surgeon's supervision.

The *human-machine interface* includes an *online display* system that combines data generated in presurgical planning with data transmitted from the robot controller to show progress of the cutting procedure superimposed on the CT-derived image views used in planning. A gas-sterilized *hand-held terminal* allows the surgeon to interact with the system during the course of the operation. This terminal supports manual guiding, motion enable, emergency power on/off, and menu selection functions. It may also be used to pace transitions from one major application step to the next and to select appropriate pre-programmed error recovery procedures should the need arise. Each of the major control components (robot control and motion checker) is able to freeze all robot motion or to turn off manipulator and cutter power in response to recognized exception conditions. If this happens, the surgeon must explicitly re-enable motion from the hand-held terminal.

V. PRESURGICAL PLANNING SYSTEM

A. Input Processing

One mundane, but nevertheless essential, task is to load the image data into the computer. The CT scanner used for the veterinary clinical trial of this system produced images on magnetic tape in GE 9800 format. The voxel size for typical scans was 0.39×0.39 mm $\times 1.5$ mm thick. Multiple cross-sectional images spaced 3 mm apart were taken throughout the proximal femur. In the vicinity of the locator pins, the images were spaced only 1.5 mm apart (i.e., they were contiguous). The input software includes facilities for tape reading, previewing image slices, selecting a region of interest to reduce the size of data sets, maintaining patient information, etc.

B. Pin Location Algorithms

A key problem is determining the location of the top center point of each locator pin relative to CT coordinates. This is by no means trivial. Although the density of the pins is much higher than that of bone, simple segmentation based on thresholding is complicated by blooming and other artifacts associated with the image formation process, so that the images are rather noisy. In particular, edge information is very unreliable.[4] The pins are not nicely aligned with the CT slices, and the CT voxels are not cubes. Even in the absence of noise, CT cross-sections that pass through the screw threads, hexagonal drive hole, and the pin head and shaft can produce images that are rather difficult to analyze. To overcome these problems, a robust three phase method has been developed.

In the first phase, simple density thresholding is used to distinguish the metallic pin voxels from surrounding "tissue"

[4]Experiments by one of the authors [18], [25] with various materials showed that titanium and ceramic yielded the best contrast without excessive blooming. However, the resulting images were still far from clean. Titanium was chosen for reasons of biocompatibility and because it is more commonly used in orthopaedic implants than is titanium.

Fig. 6. Operating room scene from first canine clinical trial in May, 1990. The surgeon is Dr. Paul. The patient was a family pet needing hip replacement surgery.

voxels. Unfortunately, blooming causes many "tissue" voxels to be mislabeled "pin," giving the pins a ragged "starburst" appearance. These artifacts are cleaned up by first dilating and then eroding the binary thresholded image with standard 3D morphology filters, using spherical structural elements. This process also smooths out the screw threads and fills in the drive socket of the pin image.

In the next phase, the approximate position and orientation of the pin are determined by calculating the first and second moments of the binary pin image.

$$m_1 = \frac{\sum_j p_j}{\sum_j 1}$$

and

$$M_2 = \frac{\sum_j (p_j - m_1)(p_j - m_1)^T}{\sum_j 1}$$

where the p_j are the coordinates of all voxels j classified "pin." Since the pin is cylindrically symmetric, two of the eigenvectors of M_2 will be practically equal. The other eigenvector, a, represents the principal axis of the pin.[5]

[5] This method would not work if the length of the pin shaft was such that all three eigenvectors had the same length. In this case, it would be necessary to use higher order moments to disambiguate the axes. However, our locator pin design precludes this possibility.
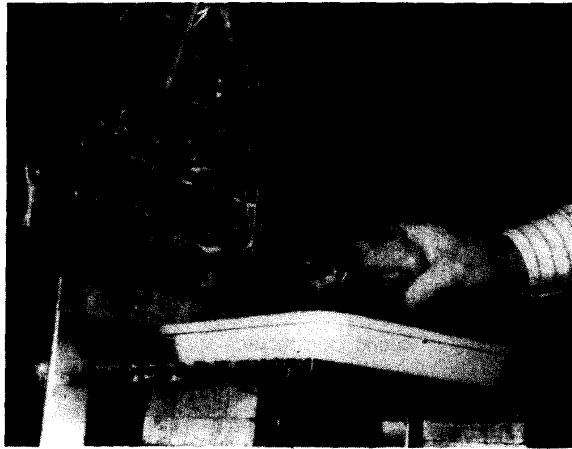
Fig. 7   Robot's wrist during shape cutting experiment. The LED beacon plates used by the motion monitoring system are clearly visible. The force sensor is just visible behind the top of the plates.

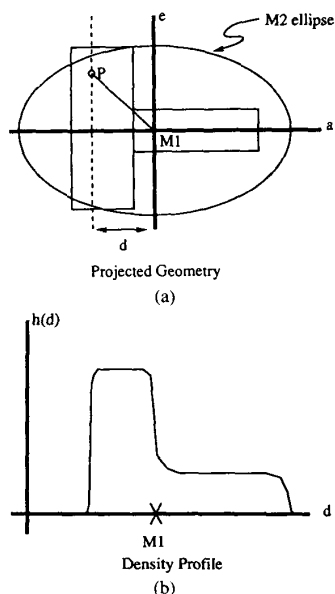
Fig. 9.   Presurgical planning display.

Fig. 8.   Projected pin profile. (a) Projected geometry. (b) Density profile.

In the third phase, a cross-sectional volume profile $h(d)$ is computed as a function of the distance d along the axis $m_1 + da$ (Fig. 8). The intercept d_0 of the "leading edge" of the pin profile is computed, and the top center point p_{tc} of the pin is then readily computed from

$$p_{tc} = m_1 + d_0 a$$

C. Interactive Docking Subsystem

The interactive docking subsystem integrates 3D image display and computer graphics techniques to support positioning of a 3D CAD model of the desired prosthesis shape relative to the CT image of the patient's anatomy. Since 3D perspective projections inherently distort distance and shape, we chose to use orthogonal 2D cross-sections to represent the 3D
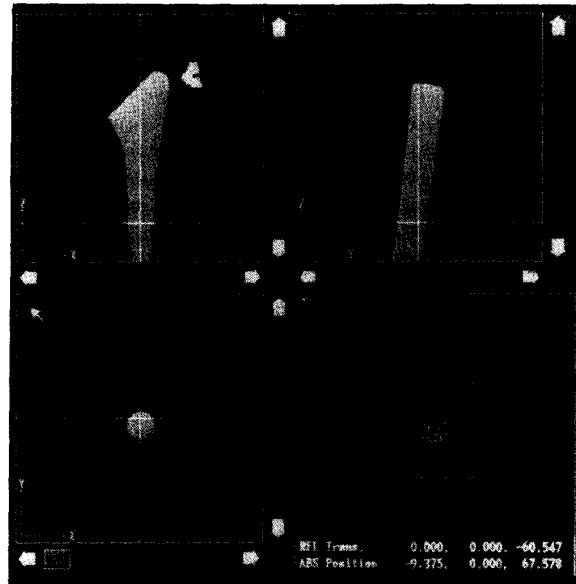
information. The interactive display screen is shown in Fig. 9. Three orthogonal sections through the CT data set representing the bone are shown, together with a simple graphic view showing the location of the three cutting planes relative to the data set. Standard resampling techniques are used to generate undistorted cross-sectional images, which may be displayed in one of three modes. Grey-scale mode simply displays the CT densities of each (resampled) voxel. Color-map mode uses different hues (red, blue, etc.) to represent different tissue classes (cortical bone, trabecular bone, etc.), which are presently computed by relatively simple intensity thresholding techniques.[6] Surface contour mode shows a graphic representation of boundaries between tissue types. This graphic data can be manipulated very quickly, and is most useful when the surgeon is identifying the desired cross-sectional views through the CT data.

In use, the surgeon typically selects boundary mode and uses the mouse to position and orient the cutting planes relative to the CT data. The surgeon then selects either grey-scale or color-map mode. Again using the mouse, the surgeon selects the desired implant model from a library of available designs, and manipulates the position and orientation of the implant relative to the CT coordinate system. As he does this, the computer automatically generates the cross-sections corresponding to the selected orthogonal cross sections and displays them superimposed on the corresponding 2D images. All manipulations, whether of the implant or of the cross-sectional CT views are specified relative to one of the three 2D views. Thus, complex 6D reorientations are accomplished by breaking them down into a sequence of simpler transfor-

[6] The threshold values used to distinguish between different bone classes were qualitatively determined by the co-author who is a surgeon (Dr. Paul), and reflect his best judgement as to what is useful. Any such distinctions are to some extent arbitrary.

mations. When the surgeon is satisfied, the coordinates of each locator pin, the implant specification, and the desired implant position and orientation relative to CT coordinates are written to a file.

In the future, we expect that the computer will assist the surgeon by computing and displaying appropriate goodness-of-fit measures and eventually proposing optimized positions and custom implant designs. Even in its present state of development, however, this system has proved to be very effective and quite easy to use. The 2D cross sectional displays are intuitively attractive to and easily learned by the orthopaedic surgeons who are the targeted end-users. The restriction to one 2D rotation or translation at a time has similarly proved to be inconsequential since our users tend to think of rotations and translations that are easily perceivable in a single display—namely the ones that the system allows on a single interaction.

VI. GEOMETRIC CALIBRATION

Geometric calibration (e.g., [26]–[28]) is a crucial component of any practical robotic application, especially one in which geometrically accurate paths are an important factor. This is equally true of surgical applications. At the same time, it is important to define methods that are simple, robust, do not require elaborate equipment, and are appropriate for the accuracies required by the task. In this section, we will describe our approach to these tradeoffs.

A. Find Pin Routine

The methods used in the calibration and in the actual surgical execution are very similar to methods earlier used in training a robot to copy pilot hole positions for automatic drilling of aircraft wing panels [29]. A ball probe cutter is inserted into the collet of the cutting tool and the force sensor is used to determine points of contact with the object being located (typically, a cylindrical pin). Points of contact are located by moving the ball to the proximity of the surface and then executing a slow guarded motion in a specified direction. As soon as the force exceeds a specified threshold, the motion is stopped. Since there may be an unpredictable amount of overshoot, a sequence of very small steps x_i are then taken in the reverse direction, and the forces f_i along the motion direction are measured at each point. The apparent compliance is estimated by a straight line approximation

$$f_i = K(x_i - x_0)$$

The point x_0 where the force goes to 0 is assumed to be the contact point. Experience has shown that this method, while somewhat tedious, is in practice very robust. Repeatabilities of the order of 25 μm are routinely obtained. A cylindrical object like a pin or cup is then easily located by locating three points on the top surface and three points on the side.

B. Kinematic Model

As stated earlier, the robot is a modified SCARA manipulator augmented by an extra pitch axis, which (in turn) carries a

six degree-of-freedom force sensor and a high speed revolute surgical cutter. The nominal kinematics are given by

$$p_{\text{tool}} = p_{\text{wrist}} + R(z, \theta_4)R(y, \theta_5)v_c$$

where $R(a, \theta)$ is a rotation by angle θ about axis a and

$$p_{\text{wrist}} = R(z, \theta_1)(l_1 x + R(z, \theta_2)l_2 x) + \theta_3 z$$

l_1 = length of first link

l_2 = length of second link

θ_1 = first joint rotation

θ_2 = second joint rotation

θ_3 = sliding joint displacement

θ_4 = roll joint rotation

θ_5 = pitch joint rotation

v_c = cutter displacement vector

There are, of course, a number of error terms corresponding to link dimensional variations, encoder offsets, etc. The calibration performed by the robot manufacturer characterizes these values quite well, and the local accuracy of the basic SCARA has proved to be sufficient for our purposes.[7] However, we were somewhat more concerned about the pitch motor and end effector, and therefore decided to develop an additional calibration procedure for these distal parts of the system. The crucial factor is the position of the tool tip, which is given by

$$p_{\text{tool}} = p_{\text{wrist}} + R(z, \theta_4 + \Delta\theta_4) \bullet (\alpha x + v_{\text{distal}})$$

where

$v_{\text{distal}} = R(x, \beta) \bullet (R(y, \theta_5 + \Delta\theta_5)(v_c + \Delta v_c)$

$\Delta\theta_4$ = rotational misalignment of joint 4 with joint 5

α = displacement of pitch axis from roll axis

β = pitch axis tilt error term

$\Delta\theta_5$ = combined pitch offset & shaft alignment error

Δv_c = cutter shaft displacement vector uncertainty

Tool orientation is relatively less important, and no special efforts were required to calibrate it, aside from determining the angular offsets $\Delta\theta_4$ and $\Delta\theta_5$.

C. Parameter Estimation

Our present calibration method uses a single vertical post rigidly mounted to robot's base. Essentially, the calibration works by repeatedly executing the "find pin" routine to measure the apparent position of the post for a number of different roll and pitch orientations. Since the post does not move, the post location and the unknown kinematic parameters α, β, $\Delta\theta_4$, $\Delta\theta_5$, and Δv_c may be found by least squares regression

---

[7] The specified repeatability of the robot we used is $\pm$ 0.05 mm in the XY plane and $\pm$ 0.02 mm in Z. The robot's specified XY region accuracy is 0.2 mm over a 250 mm square. Over the rather shorter distances involved in machining a canine implant, the accuracy rapidly approaches the repeatability, which, in our experience, was actually better than the specified value.

on the linearized relation

$$p_{\text{post}} \cong p_{\text{wrist}} + R_4 \bullet R_5 \bullet (v_c + \Delta v_c)$$
$$+ \beta R_4 \bullet (x \times R_5 \bullet v_c)$$
$$+ \Delta \theta_4 R_4 \bullet (z \times R_5 \bullet v_c)$$
$$+ \Delta \theta_5 (y \times (R_4 \bullet R_5 \bullet v_c)) + \alpha R_4 x$$

where

$$R_4 = R(z, \theta_4)$$
$$R_5 = R(y, \theta_5)$$

Our experience with this calibration procedure has been quite good. A typical calibration run consisted of 28 poses, with roll angles varying through $\pm 90°$ and pitch angles varying from $2°$ to $60°$. After reduction of the data, the average residual variation in the apparent position of p_{post} was typically about 0.1 mm. Over a series of 9 calibration runs made during the canine clinical trial, the average residual variation ranged from 0.08 mm to 0.12 mm, and the maximum residual magnitude ranged from 0.16 mm to 0.33 mm, which is well within the required accuracy for this application. Further, since the wrist orientation does not change during the shape cutting phase, any remaining wrist calibration error simply causes the position of the hole to be shifted slightly in the patient's femur, and does not affect the actual shape being cut.

VII. SHAPE CUTTING

At present, we use a constant orientation cutting strategy. The shape is cut as follows:

1) An end mill (typically, 7–9mm in diameter) is placed in the collet of the cutting tool.
2) The cutter is oriented parallel to the long axis of the implant.
3) Successive transverse pockets (typically, about 2.5 mm deep) are cut to produce the rough shape of the implant. At the conclusion of this stage, the implant shape has a stair case appearance.
4) Successive longitudinal cuts are made to remove the excess material in the stair cases. Although this produces a slightly scalloped surface finish, in practice it is easy to approximate the desired surface with a relatively small number of cuts. The residual height δh of any scallop will be given by

$$\delta h = r_{\text{cutter}} - \sqrt{r_{\text{cutter}}^2 - \frac{d^2}{2}}$$

where r_{cutter} is the cutter radius and d is the distance between cuts. Solving for d gives

$$d = \sqrt{4 r_{\text{cutter}} \delta h - 2(\delta h)^2}$$
$$\cong 2\sqrt{r_{\text{cutter}} \delta h} \quad \text{for small } \delta h$$

Thus, $r_{\text{cutter}} = 5$ mm and $\delta h = 0.05$ mm would require the finishing cuts to be 1 mm apart. These cuts are made without changing the cutter orientation.

5) If necessary, the end mill is replaced with a smaller diameter ball cutter and additional finishing cuts are

made to sharpen the corners of the implant hole. Obtaining proper clearance for these cuts requires the cutter orientation to be changed slightly. Changes in the implant designs during veterinary clinical testing rendered this step unnecessary.

One advantage of a constant orientation cutting strategy is that it substantially eliminates the effect of unmodeled kinematic errors in the distal parts of the robot on the shape of the hole being cut, although they do still affect the location of the hole relative to the bone. Since the shape dimensional tolerances are in fact somewhat tighter than the positioning tolerances, maintaining a constant orientation is indeed valuable. More complex implant shapes, of course, will require full five axis trajectories.

VIII. INTRAOPERATIVE DISPLAY

The presurgical planning system is also used in the operating room to provide displays showing the progress of the cutting phase of the surgery. During surgery, the planning system is connected to the robot controller via a standard serial communication line, and rechristened the real time monitor. Three orthogonal cross-sections through the 3D CT data set used to plan the surgery are displayed together with corresponding cross sections of the shape to be cut, just as in presurgical planning. As each successive cutting stroke is made, the robot controller sends short messages to the display computer, which then changes the color of the portions of the cross-sectional images corresponding to the cutting stroke. Once a complete layer is cut out, that entire portion changes color yet again.

IX. SAFETY CHECKING SUBSYSTEMS

A. Requirements

Safety was a primary consideration in designing the system. The principal requirements were defined by the co-authors of this paper who are surgeons (Dr. Paul and Dr. Bargar). These included:

1) *The robot must never run away.* No single-mode hardware (or system) error may cause the application software to lose control of its motions. Furthermore, the application software must request only proper motions.
2) *The robot must never exert excessive force on the patient.* If forces on the cutter exceed expected values by more than a predefined threshold amount, then something may be wrong, and the robot must stop moving immediately.
3) *The robot's cutter must stay within a prespecified positional envelope relative to the volume being cut.* For hip replacement surgery, the main goal is to prevent a systematic shift in the placement or shape of the hole; a single gouge is generally reparable, although undesirable. Of course, other surgical procedures (like brain surgery) may be less forgiving.
4) *The surgeon must be in charge at all times.* This is, of course, the fundamental dilemma. The surgeon has to trust the system to some extent. Nevertheless, the system must provide the surgeon with timely information

about its status, and the surgeon must be able to pause motion at any time. Once robot motion is stopped, he or she must be able to further query the robot's status, to manually guide it, to select an appropriate recovery procedure to continue the surgery, or to completely terminate use of the robot and continue manually.

B. Robot Controller Checks

The robot controller routinely performs many safety and consistency checks, including monitoring position and velocity limits in the joint servos and monitoring of external signals. In addition to a basic power-enable relay (external to the controller), controller software provides facilities for disabling manipulator power, for freezing or pausing motion, for resuming interrupted motions, and for transferring control to application software recovery procedures. A safety time-out monitor turns off arm power if the controller does not affirmatively verify system integrity every 18 ms.

Many conditions (externally signalled consistency checks, force thresholds, pushbutton closures, etc.) interrupt the application program, pause motion, or drop power under certain conditions. The surgeon can then use the hand-held terminal to query system status, to select local actions (such as manual guiding or withdrawal of the cutting tool), to continue the present motion, to discontinue or repeat the present step of the procedure, or to restart from an earlier stage of the procedure. One very common case is a simple surgeon-initiated pause to allow the surgical team to perform some housekeeping function like replacing an irrigation bottle or to allow the surgeon to satisfy himself or herself that all is well.

C. Force Monitor Checks

The microprocessor interface to a wrist-mounted force sensor computes forces and torques at the cutter tip. If any tip force component greater than approximately 1.5 kgf is detected, the controller is signalled to pause motion. Forces greater than about 3 kgf cause arm power to be dropped. Experiments in which a sudden large motion is commanded in the middle of cutting confirm that these checks are quite effective in detecting run away conditions. They are also effective in detecting such conditions as the cutter stalling or being impeded by improperly retracted soft tissue.

D. Independent Motion Monitoring Checks

We developed an independent checking subsystem to verify that the cutter step stays within a defined "safe" volume relative to the bone, essentially corresponding to the implant shape and an approach region. The checking system is implemented on a separate PC/AT computer from the robot controller, in order to minimize the chances of common mode failures. The check requires two steps: 1) verification that the bone does not move relative to the fixator, which is rigidly attached to the robot's base, and 2) verification that the end effector never strays from a defined volume in space.

We devised a strain gauge system for detecting motions of the bone relative to the fixator. Bench experiments demonstrated that motions on the order of 0.1 mm could be detected. However, experiments with the fixator indicated that even

rather large forces (5 kgf) produced only negligible (16 μm) motion, and the bone motion monitor was not used in any clinical tests.

To verify end effector motion, we used a Northern Digital Optotrak^{TM} 3D digitizer, which is capable of tracking light emitting diodes to an accuracy of better than 0.1 mm at a rate of approximately 1000 positions/second. We fabricated a rigid PC card with eight such beacons and affixed it to the robot's wrist, as shown in Fig. 7. An arbitrary coordinate system for the PC card was defined from the beacon positions, and the positions b_π of the beacons relative to this coordinate system were measured. The Optotrak measures the positions b_{oi} of these beacons in space, and computes a best estimate of the plate position F_p by regression from the relationship

$$b_o \cong F_p \bullet b_p$$

The robot-to-Optotrak and plate-to-cutter transformations, T_{ro} and T_{pc} are computed by ordinary least squares estimation from data taken with the robot in various known positions, using appropriate linearized models. Using these transformations, an estimate of the cutter coordinates F_{rc} relative to the robot may be obtained from the relationship

$$F_{rc} = T_{ro}^{-1} \bullet F_p \bullet T_{pc}$$

Constructive solid geometry (CSG) tree "check volumes" corresponding to implant and cutter selection were constructed from primitives bounded by quadric surfaces

$$p_c^T \bullet Q_i \bullet p_c + q_i \bullet p_c + d_i \leq 0$$

located one millimeter outside the furthest nominal excursions of the cutter when the shape is cut. Intraoperative checking is performed by reading the beacon plate coordinates from the Optotrak, computing the corresponding cutter position, and then checking to see if this position falls outside the check volume. If so, the checking subsystem signals an out-of-bounds condition through an optically isolated digital port to the robot controller, which pauses motion and then obtains more detailed information through a serial communications line.

To verify the performance of this system, we deliberately moved the cutter in a succession of very small steps through the boundary of the checking volume. We found that the system could detect when a motion crossed a threshold to approximately 0.2 mm precision with constant orientation, and approximately 0.4 mm with cutter reorientation. Checking rates of approximately 3–4 Hz were obtained using a slow (6 MHz '286) PC/AT. At typical cutter speeds, the total excursion before motion is frozen is about 2 mm after accounting for all latencies.

X. EXPERIENCE AND DISCUSSION

A. In Vitro

Extensive tests were conducted on plastic and cadaver bones and on foam test blocks, in order to verify basic system accuracy and to gain confidence in overall system behavior [25]. Fig. 11. shows typical cross sections produced
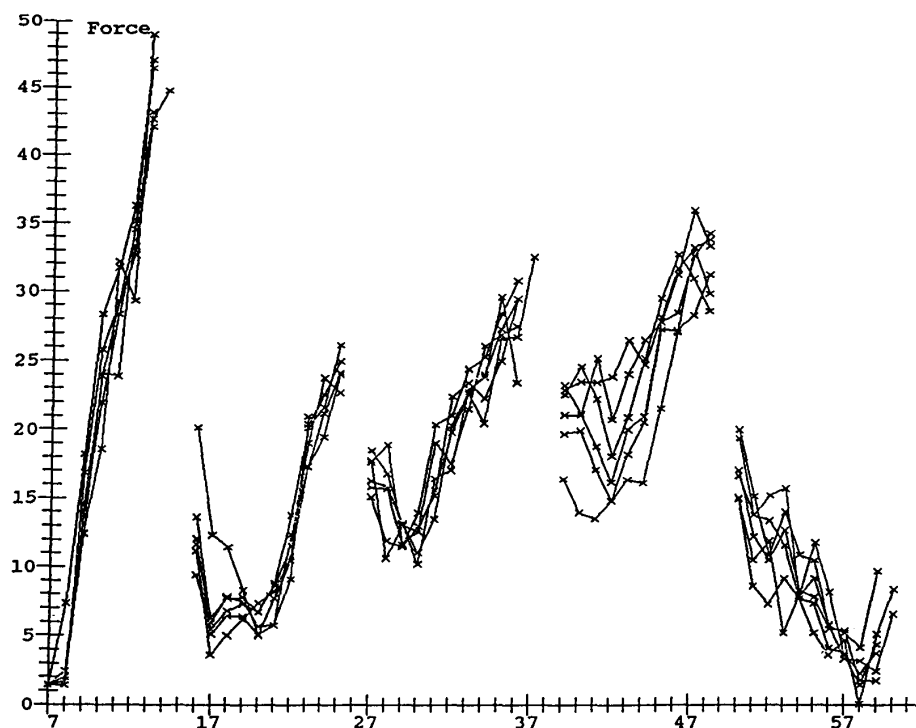
Fig. 10. Cutting forces on cadaver dog bone. This figure shows a plot of the magnitude of the forces on the cutter tip for five typical cross-sectional cuts. In this plot, 40 force units correspond to one pound (0.45 kgf).

by manual broaching and robotic machining. In one bottom line experiment, three pins were implanted into a test fixture and located on CT images of the fixture. A number of foam blocks were then successively (and repeatably) clamped into a socket in the test fixture, which was placed at various poses within the workspace of the robot. Test shapes were cut in the foam blocks, using steps 3 through 5 of the surgical procedure. The positions of these shapes and of the pins were then measured on a coordinate measuring machine with an accuracy of approximately 0.0125 mm (0.0005 inches). In a typical test run the three blocks were cut for each of four separate poses (all combinations of "left leg/right leg" and "0 degrees/15 degrees" fixator pitch) for each of three separate CT scans, for a total of $3 \times 4 \times 3 = 36$ blocks. The total placement error of the test shapes was found to be 0.5 mm for a test fixture with pins placed at the distances they would be on a human. Similarly, the dimensions of test shapes machined in cadaver bone were measured with calipers accurate to 0.02 mm. Dimensional errors were less than 0.05 mm. Further tests, in which actual implant shapes were cut in both foam blocks and canine cadaver bones and an implant was then inserted into the hole, were also conducted. Although a dimensional study similar to [17] was not performed, the fit achieved was qualitatively very good. The implant slipped into the hole with little effort and fit snugly. No gaps could be seen if the foam block was split longitudinally after insertion of the implant.

Cutting forces for bone machining were also measured. Typical results are shown in Fig. 10. The greatest force, which

seldom exceeded 0.5 kgf, typically was encountered when the cutter moved from the center "plunge" position to the first corner of a rectangular section of bone being removed. Forces would then drop off substantially as the cutter began moving along edges of the section, building up to somewhat smaller local maxima (about 0.3–0.4 kgf) as successive corners of the shape were reached.

Except for bone motion detection, the redundant checking mechanisms discussed above were all integrated into the prototype surgical system and used in in vitro testing on cadaver bones. The most common error condition detected during these tests was excessive cutting force when the cutter plunged into unusually hard bone, which caused the controller to pause robot motion. Continuation from this condition was easily achieved by backing up the cutter 1–2 mm and restarting the current cut.

Both the real time monitor and the motion tracking system proved to be surprisingly useful in application debugging. Even though the display was essentially an animation, it provided useful information about exactly where the robot should be and what the controller thought it was doing. The motion tracker provided a useful consistency check to the calibration procedures. It also caught a real bug in the shape cutting code that might otherwise have been very hard to find.

B. In Vivo

A clinical trial on 26 dogs needing hip replacement surgery was conducted from May 1990 through September 1991. All
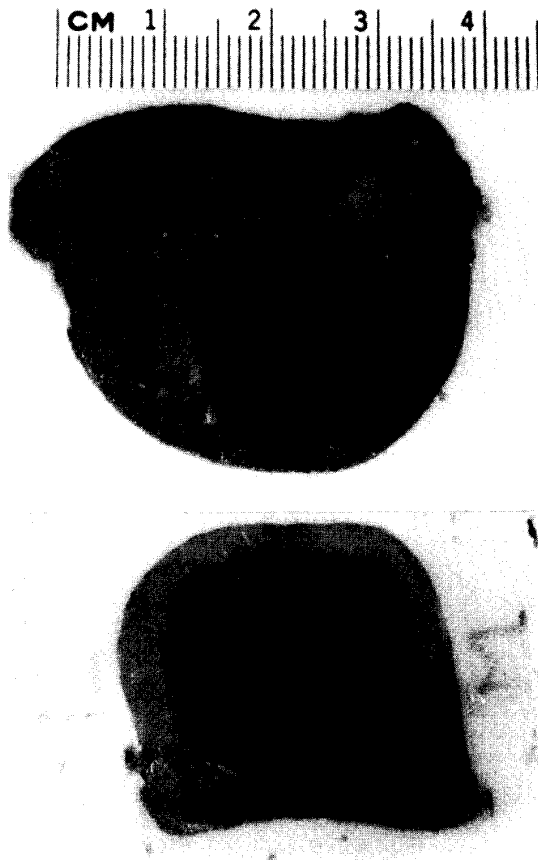
Fig. 11. Comparative cross section. This figure shows sections of human cadaver bone (top) prepared with a manual broach and (bottom) machined with a robot. The measured midline dimensions of the machined sections are within 0.05 to 0.1 mm of nominal. Some surface irregularities are seen where the machined surface intersects chancellous bone; other chipping arises from the bone-sectioning process.

procedures were successful, with no intraoperative complications or infections. There were no intraoperative or postoperative cracks or fractures, and (in the opinion of the surgeon) the implants were easily inserted and provided more mechanical stability than would normally be experienced with manual broaching of the femur. In contrast, cracking[8] was experienced in 5 out of 15 cases in a manually broached control group, and the implant placement was not as good. Radiographs were used to compare placement of the implant in the femur for 15 cases in which the robot was used to prepare the femur with 15 cases in which conventional manual broaching was done. It was found that, for the Techmedica cementless canine implants used in the study, conventional broaching often resulted in the proximal end of the implant was often tilted more toward the medial direction of the femur (i.e., more in a varus orientation) than the surgeon judged to be optimal. Because of this possibility, the surgeon tended to select a slightly smaller implant design than he otherwise would have. In contrast, robotic machining consistently matched the implant axis with

[8] Fixed by wrapping cerclage wire around the affected bone.

the axis of the proximal femur, enabling the surgeon to select the implant size that best matched the patient's internal bone geometry.

In surgery, the system worked very well. Although no systematic effort was made to compare surgical execution times, the total time of surgery was roughly comparable to that for manual broaching. Anecdotally, it was observed that the time required for robot machining was more consistent from case to case than the time required for manual broaching. There were very few glitches, and that there was little actual use of any of the error recovery capabilities of the system. This is as it should be. The force monitor occasionally froze motion when the cutter encountered an unusually hard section of cortical bone at the proximal end of the femur. In these cases, the surgeon simply restarted motion with the hand-held terminal. On two other occasions where the force monitor stopped motion (once when the cutter became entangled in some suture material and once when it got caught in an assistant's glove) it was necessary for the robot to withdraw from the bone. The surgical team cleared the entanglement and resumed the procedure with only a few cutting motions being repeated.

The veterinary surgeon (Dr. Paul) relied on the real time monitor to provide positional status information, in conjunction with his other senses. By listening to the cutter, he could tell when the cutter was in contact with hard bone. When he heard a change in pitch, he would look at display to verify that what he was hearing was consistent with where the robot was cutting. One interesting possibility for future work would be to automate such multisensory cross-checking.

The separate motion checking system was not used *in vivo* for dogs. The surgical field is rather crowded, since a technician must constantly irrigate the bone while the robot is cutting it. It proved to be very difficult to place the vision system sensors in the veterinarian's operating room so that they would always have a clear view of the end effector. One possibility would have been to mount the cameras overhead. Another would have been to use the system for occasional *spot checks* of the robot. However, one consequence of the confidence gained from the *in vitro* tests was that the surgeon concluded that the additional redundancy gained was not worth the added complexity for veterinary cases. A different sensing solution altogether, based on monitoring of redundant joint encoders, was consequently adopted for the subsequent *human-qualified* second-generation system [19]–[22].

XI. CONCLUSIONS

The system described in this paper demonstrated the feasibility of adapting a general-purpose manipulator for use as a precise surgical tool. We were been able to demonstrate an order-of-magnitude improvement in the precision with which a surgeon can execute a critical step in hip replacement surgery.

Beyond this, it may be worthwhile to recap how this system has addressed the general requirements discussed at the beginning of the paper. We found that even very simple human-machine interaction technology can be surprisingly effective, although further improvements are desirable. The use

of hands-on, force compliant guiding for positioning the robot has been especially successful, since it enables the surgeon to position the robot in a way that is natural and intuitively simple. An animated information display showing the progress of the surgical procedure was useful, although this is again an area where considerable improvement can be made both in the presentation of information to the surgeon and in the incorporation of realtime sensing. One ultimate system might be some sort of heads-up display showing the surgical plan superimposed on the actual patient, with the display being updated based on a combination of position tracking, cutter force data, acoustic sensing, and intraoperative imaging.

Model-reality registration was accomplished in this case by the use of landmark pins, which could be located easily in both CT images and in physical reality. Although their use in this particular surgery is acceptable, less invasive methods may often be desirable. One obvious choice is to register intraoperative radiographs to features on CT-derived models. Another would be to use a 3D digitizer such as the Optotrak to point out anatomical features, and then to provide realtime tracking of markers placed on the patient at the time of surgery (see, e.g., [30]).

Verification of robot performance and of the methods chosen to register the plan to reality was an important issue, one that required as much time and effort as any other aspect of the system development, and is discussed more fully in [25]. In this regard, experimental measurement of individual error sources and the bottom line experiments described above went hand-in-hand. Similarly, the optical endpoint check, though ultimately not used in the operating room, proved very useful in debugging the shape cutting software.

The related issue of safety was also paramount. The fact that our application required a robot to move a tool in contact with a patient motivated us to implement a number of redundant consistency checking mechanisms, which proved quite valuable both in application debugging and in actual surgery.

In implementing these redundant checking mechanisms, we encountered an important tradeoff with the realities of operating room compatibility for a complex piece of equipment. The area around the patient is crowded, and it is often awkward to maintain a clear field of view required for optical checking equipment. This led the surgeon in this particular application to conclude that whatever extra safety may be gained by a completely independent visual check, compared to checks on the robot's encoders, does not justify the extra system and operating-room complexity involved.[9] In other applications where an optical system is also taking a more active role, for example, in tracking the patient's anatomy, it may be desirable to permit its use for redundant safety checking as well.

In any case, it is clear that the system reported here represents only a step in the evolution of a man-machine partnership in the operating room, in which the complementary abilities of robotic devices and humans are exploited under the human's supervision to help provide a better result for the patient. Indeed, this process has continued for the hip

[9] As mentioned earlier, the next generation robot [22] incorporates an additional, independent set of encoders to provide further redundancy.

surgery augmentation system that we have described. The experience gained with veterinary patients provided the basis for development of a second generation system [19]–[21] for use on humans.

REFERENCES

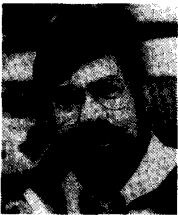[1] Y. S. Kwoh, J. Hou, E. Jonckheere, and S. Hayati, "A robot with improved absolute positioning accuracy for CT guided stereotactic surgery," *IEEE Trans. Biomed. Eng.*, pp. 153–161, February 1988.

[2] S. Lavallee, "Gestes medico-chirurgicaux assistes par ordinateur: Application a la neurochirurgue stereotaxique," Ph.D. thesis, University of Grenoble, 1989.

[3] S. Lavallee, "A new system for computer assisted neurosurgery," in *Proc. 11th IEEE Engineering in Medicine and Biology Conf.*, Seattle, Nov 1989, pp. 926–927.

[4] Bruce A. Kall, Patrick J. Kelly, and Stephan J. Goerss, "Interactive stereotactic surgical system for the removal of intracranial tumors utilizing CO_2 laser and CT-derived database," *IEEE Trans. Biomed. Eng.*, pp. 112–116, February 1985.

[5] Patrick J. Kelly, Bruce A. Kall, Stephan Goerss, and Franklin Earnest, "Computer-assisted stereotaxic laser resection of intra-axial brain neoplasms," *J. Neurosurg*, pp. 427–439, March 1986.

[6] James M. Drake, Michael Joy, Andrew Goldenberg, and David Kreindler, "Robotic and computer assisted resection of brain tumors," in *Proc. Fifth Int. Conf. on Advanced Robotics*, Pisa, June 1991, pp. 888–892.

[7] D. Glauser, P. Flury, N. Villotte, and C. W. Burckhardt, "Conception of a robot dedicated to neurosurgical operations," in *Proc. Fifth Int. Conf. on Advanced Robotics*, Pisa, June 1991, pp. 899–904.

[8] B. L. Davies, R. D. Hibberd, A. Timoney, and J. Wickham, "A surgeon robot for prostatectomies," in *Proc. Fifth Int. Conf. Advanced Robotics*, Pisa, June 1991, pp. 871–875.

[9] K. Ikuta, M. Tsukamoto, and S. Hirose, "Shape memory alloy servo actuator system with electric resistance feedback and application for active endoscopes," in *IEEE Robotics and Automat. Conf.*, 1988.

[10] R. H. Sturges and S. Laowattana, "A flexible tendon-controlled device for endoscopy," in *Proc IEEE Robotics and Automation Conference*, Sacramento, CA, 1991.

[11] Philip Green, "Advanced teleoperator technology for enhanced minimally invasive surgery," in *Proc. Medicine Meets Virtual Reality Conference*, San Diego, June 4–7 1992.

[12] Yulan Wang, "Robotically enhanced surgery," in *Proc. Medcine Meets Virtual Reality II*, San Diego, Jan 27–30 1994.

[13] J. A. McEwen, C. R. Bussani, G. F. Auchinleck, and M. J. Breault, "Development and initial clinical evaluation of pre-robotic and robotic retraction systems for surgery," in *Proc. Second Workshop on Medical and Health Care Robots*, pp. 91–101, Newcastle-on-Tyne, Sept. 1989.

[14] J. L. Garbini, R. G. Kaiura, J. A. Sidles, R. V. Larson, and F. A. Matson, "Robotic instrumentation in total knee arthroplasty," in *Proc. 33rd Annual Meeting, Orthopaedic Research Society*, San Francisco, January 1987, p. 413.

[15] Thomas C. Kienzle, S. David Stulberg, Michael Peshkin, Arthur Quaid, and Chi-haur Wu, "An integrated CAD-robotics system for total knee replacement surgery," in *Proc. 1993 IEEE Conf. Robotics Automat.*, Atlanta, May 1993, pp. 889–894.

[16] P. M. Sandborn, S. D. Cook, W. Spies, and M. Koster, "Tissue response in porous coated implants locking initial bone apposition," *J. of Arthroplasty*, December 1988.

[17] H. A. Paul, D. E. Hayes, W. L. Bargar, and B. D. Mittelstadt, "Accuracy of canal preparation in total hip replacement surgery using custom

broaches," in *Proc. First International Symposium on Custom Made Prostheses*, Dusselforf, October 1988.

[18] H. A. Paul, B. D. Mittelstadt, B. L. Musits, W. L. Bargar, and D. E. Hayes, "Application of CT and robotic technology to hip replacement surgery," in *Proc. First International Symposium on Custom Made Prostheses*, Dusseldorf, October 1988.

[19] H. Paul, B. Mittelstadt, W. Bargar, B. Musits, R. Taylor, P. Kazanzides, J. Zuhars, B. Williamson, and W. Hanson, "A surgical robot for total hip replacement surgery," in *Proc. IEEE Conference on Robotics and Automation*, Nice, May 1992.

[20] P. Kazanzides, J. Zuhars, B. Mittelstadt, P. Cain, F. Smith, L. Rose, and B. Musits, "Architecture of a surgical robot," in *Proc. 1992 IEEE Conf. on Syst. Man Cyber.*, Chicago, August 1992.

[21] Brent Mittelstadt, Peter Kazanzides, Joel Zuhars, Phil Cain, and Bill Williamson, "Robotic surgery: Achieving predictable results in an unpredictable environment," in *Proc. Sixth Int. Conf. on Advanced Robotics*, Tokyo, Nov 1–2 1993.

[22] B. D. Mittelstadt, P. Kazanzides, J. Zuhars, B. Williamson, P. Cain, F. Smith, and W. Bargar, "The evolution of a surgical robot from prototype to human clinical use," in *Computer Integrated Surgery*, R. Taylor, S. Lavallee, G. Burdea, and R. Moesges, Eds. Cambridge, MA: MIT Press. To appear.

[23] W. H. Hanson, H. A. Paul, Bill Williamson, and Brent Mittelstadt, "Orthodock: A Computer System for Presurgical Planning," in *Proc. 12th IEEE Medicine & Biology Conf.*, Phila., 1990.

[24] Russell H. Taylor, Peter Kazanzides, Brent D. Mittelstadt, and Howard. A. Paul, "Redundant consistency checking in a precise surgical robot," in *Proc. 12th IEEE Medicine & Biology Conf.*, Phila., November 1990.

[25] B. D. Mittelstadt, Ph.D. thesis, University of California at Davis, in preparation.

[26] John M. Hollerbach, "Advances in robot calibration," in *Proc. Sixth ISRR*. Cambridge, MA: MIT Press, 1993.

[27] L. J. Everett, "Models for diagnosing robot error sources," in *Proc. 1993 IEEE Conf. on Robotics and Automation*, Atlanta, May 1993, pp. II.155–159.

[28] A. Goswami, A. Quaid, and M. Peshkin, "Complete parameter identification of a robot from partial pose information," in *Proc. 1993 IEEE Conf. on Robotics and Automation*, Atlanta, May 1993, pp. I.168–173.

[29] Russell H. Taylor, "Method and apparatus to teach a robot the position and orientation of hole centerlines." US Patent 4,4485,453, filed 1982, issued Nov 27, 1985.

[30] Russell H. Taylor, Court B. Cutting, Yong-yil Kim, Alan D. Kalvin, David Larose, Betsy Haddad, Deljou Khoramabadi, Marilyn Noz, Robert Olyha, Nils Bruun, and Dieter Grimm, "A model-based optimal planning and execution system with active sensing and passive manipulation for augmentation of human precision in computer-integrated surgery," in *Proc. 1991 Int Symposium on Experimental Robotics*, Toulouse, France, June 25–27 1991. New York: Springer-Verlag, 1991.

Brent Mittelstadt received a degree in physiology from the University of Arizona, and expects to receive the Ph.D. degree in biomedical engineering from the University of California at Davis in 1994. He is currently a biomedical engineer, and began working on the ROBODOC project as an IBM research fellow. He is responsible for much of the development work on the robot arm.

Howard A. Paul, deceased, received the Veterinary Degree from the Ecole Nationale Veterinaire, D'Alfort, France, and completed his residency at the University of California at Davis School of Veterinary Medicine. Dr. Paul was the co-founder and Chairman of Integrated Surgical Systems. He had extensive research experience in orthopaedics and implant technology. He held patents on several orthopaedic devices, including a cementless hip implant.

William A. Hanson (S'74–M'85) received the B.S. and M.S. degrees in electrical engineering from Rensselaer Polytechnic Institute. In 1976 he joined the Federal Systems Division of IBM in Oswego, New York, working with advanced, real time signal processing systems. From 1985 to 1992 he was at the IBM Palo Alto Scientific Center, where he was initially involved in designing a language and simulation environment for the study of artificial neural networks. Subsequently, he worked in the image science area of numerically intensive computing. Since 1992, he has been leading the high-performance computing effort at Loral Federal Systems Corporation in Gaithersburg, MD.

Mr. Hanson's research interests include high performance computing, information theory, digital image processing (particularly three-dimensional medical image processing), visualization, and artificial neural networks. He is an IBM Certified Architect and a member of the IEEE. Photograph not available at the time of publication.

Peter Kazanzides (S'82–M'88) received the Ph.D. degree in electrical engineering from Brown University. He joined the ROBODOC project as a post-doctoral fellow in the IBM Robotics Research Laboratory and is responsible for the development of base systems. His research interests are in robotics control and programming.

Russell H. Taylor (S'68–M'77–SM'87–F'94) received the B.E.S. degree from Johns Hopkins University in 1970 and a Ph.D. in Computer Science from Stanford in 1976. He joined IBM Research in 1976, where he developed the AML language. Following a two-year assignment in Boca Raton, he managed robotics research activities at IBM Research from 1982 until returning to full-time technical work in late 1988. Since March 1990, he has been manager of Computer Assisted Surgery.

His research interests include robot systems, programming languages, model- based planning, and (most recently) the use of imaging, model-based planning, and robotic systems to augment human performance in surgical procedures.

Dr. Taylor is Editor Emeritus of the *IEEE Transactions on Robotics and Automation*, a Fellow of the IEEE, a member of the IEEE Robotics and Automation Society AdCom, co-chair of the IEEE Robotics and Automation Society Technical Committee on Medical Robotics, and a member of various other honorary societies, panels, program committees, and advisory boards.

Joel Zuhars (M'92) received the M.S. degree in electrical and computer engineering from the University of California at Davis. He was previously employed by the Johns Hopkins University Applied Physics Lab. He is currently a Lead Engineer for Integrated Surgical Systems, Inc., with the responsibility of developing the surgeon-to-robot interface.

Bill Williamson is a Senior Software Engineer for Integrated Surgical Systems, Inc. He developed the user interface for the preoperative planning workstation. He is responsible for researching and developing new imaging applications.

Edward Glassman received the B.S. degree in biomechanics from the Massachusetts Institute of Technology in 1982, where he developed a microprocessor controlled, adaptive, above-knee prosthesis. From 1982 through 1990 he was a researcher at the IBM T. J. Watson Research Center, where he focused on the development and application of high-performance, sensor-based robotic systems for a wide variety of tasks, ranging from semiconductor manufacturing to robotic surgical assistants.

Currently, he is a member of the IBM Consulting Group, specializing in Healthcare. Particular areas of interest continue to be the assessment and application of new technologies that fundamentally change the way a business operates.

Bela Musits received the B.S. and M.S. degrees in engineering and the M.B.A. degree in marketing from Rensselaer Polytechnic Institute. He is currently the President of Integrated Surgical Systems, Inc. Prior to joining Integrated Surgical Systems, Inc., he managed the ROBODOC project within IBM, and collaborated with Dr. Paul and Dr. Bargar.

William L. Bargar received the M.D. degree from Ohio State University. He completed his residency at Case Western Reserve University and his fellowship in joint replacement at UCLA Medical Center. Currently he has a private practice at Sutter General Hospital in Sacramento, CA. Dr. Bargar is a specialist in joint reconstructive surgery, and was the first surgeon to perform ROBODOC-assisted surgery on a human patient.