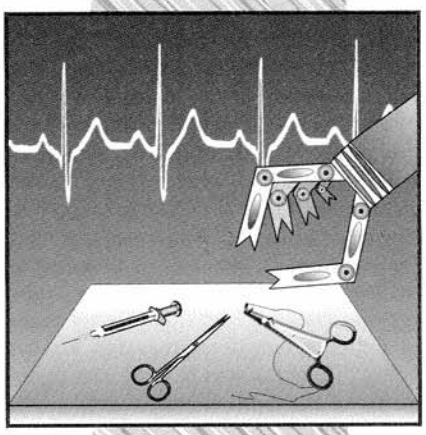

Total Knee Replacement

Computer-assisted surgical system uses a calibrated robot

Thomas C. Kienzle III[1], S. David Stulberg[1], Michael Peshkin[2], Arthur Quaid[2], Jon Lea[2], Ambarish Goswami[2], and Chi-haur Wu[3]

[1]Department of Orthopaedic Surgery
[2]Department of Mechanical Engineering
[3]Department of Electrical Engineering and Computer Science
Northwestern University

Every year, thousands of patients suffering joint disabilities such as rheumatoid arthritis or osteoarthritis undergo total knee replacement (TKR) surgery in order to return to a more active and pain-free life. Currently, in order to implant prosthetic knee components (Figs. 1 and 2), a complex jig system of cutting blocks, alignment rods, etc., is used to help the surgeon approximate the geometry of the bones and select the appropriate size and location of the components. This process, which relies heavily on an individual surgeon's experience with a given jig system, has prompted the search for a more accurate and repeatable system for the placement of total knee prosthetic components.

The aim of total knee replacement surgery is to replace the articular surfaces of the knee. Specifically, the end of the femur (thigh-bone) is replaced with a chrome-plated titanium component, and the top of the tibia (shinbone) is replaced with a polyethylene-topped titanium device. The patella (kneecap) is also resurfaced with a polyethylene component.

Total Knee Replacement Surgery

Conventional jig-based systems

Before conventional knee replacement surgery, a standard x-ray of the whole leg (front-view) is examined to determine the proper angle of the femoral component with respect to the shaft of the femur. This angle (usually about seven degrees) is chosen such that the tibia will be perpendicular to the ground and be directly under the hip joint.

During surgery, a hole is drilled at the end of the femur and a rod is placed down the center of the bone. A jig is placed on the rod, adjusted to the preoperatively determined angle, and holes are drilled into the bone where indicated by the jig. Guide pins are inserted in these holes, a cutting block is placed on the pins, and a cut is made with a powered oscillating bone saw, which defines the horizontal plane of the femoral surface. A second jig is inserted on the rod, and the femoral component's remaining cut locations are determined largely by inspection. The tibial component is placed in a similar manner, except that an alignment jig external to the leg is used to direct its positioning.

After all cuts are made, the prosthetic components are tested in place and a polyethylene spacer is chosen to maintain the proper ligament tension and full range of joint motion. The components are then cemented in place.

1. Prosthetic components used in total knee replacement surgery.

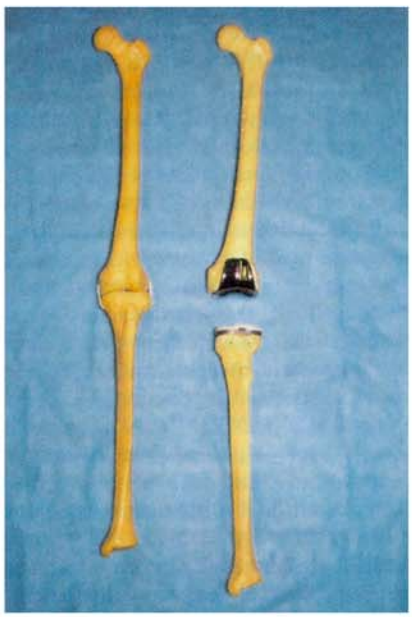

2. Femur and tibia before and after total knee replacement.

0739-5175/95/$4.00©1995

Limitations of jig-based systems

It is currently believed that the alignment accuracy of the prosthetic components affects the surgical outcome for the patient and, possibly, the longevity of implant. Using the existing TKR jig systems, components are oriented within two or three degrees of the desired "natural" position. These jig systems introduce several sources of inaccuracy in alignment of the prosthetic components. One major source of error is that only the very ends of the involved bones are exposed during surgery, forcing the surgeon to make decisions regarding bone and joint alignment based on very limited information. The preoperative x-ray can help, but still represents only a two-dimensional projection of complex three-dimensional structures.

A second source of error is the jig system itself, which represents a physical embodiment of a component placement algorithm favored by the system designer. Optimal placement of components may not be achieved when the configuration of an individual patient's bones differ from those of the generalized model assumed by the jig system, or if the algorithm on which the system is based is sub-optimal or obsolete.

Further, the existing jig systems, by necessity, direct a set of cuts in the bones based largely on local topography. It is hoped that these cuts will lead to the proper placement of the components. A preferable approach would be to visualize the correct placement of the prosthesis, based on the overall geometry of the leg, and then determine the proper cuts required to achieve optimal placement.

Motivation for robotic systems in total knee replacement

An integrated system has been developed that uses a workstation (which displays a three dimensional model of the patient's bones obtained from a CT scan of the leg), and a modified industrial robot to direct the placement of prosthetic components. The single focus of our computer-assisted TKR system is the accurate overall positioning of components. In contrast, in previous work with computer/robot total joint systems, such as the Integrated Surgical Systems ROBODOC™ total hip replacement system [1], accurate machining of local surfaces has been of vital importance. In our TKR system, a graphics computer allows prosthesis placement decisions to be made by the surgeon based on a full view of the bones involved. A component placement algorithm can be implemented in software, which can be easily altered to accommodate an individual patient's bones, or updated as better algorithms are developed. Finally, the intended component placement can be visualized and, if necessary, corrected well before any live bones are cut.

In addition to benefits from improved component placement, a computer/robotic system may ultimately allow for a smaller incision in the patient and require less time for surgery, which may both decrease complications due to infection and reduce the cost of surgery. Also, by using a computer to plan component placement and a robot to perform it, decisions made during the process are repeatable, can be accurately implemented, and are readily available for systematic study of optimal component placement geometry.

System operation

To describe the computer/robotic system developed, it is easiest to step sequentially through the TKR procedure. The key steps are:

Preoperative procedures

- Place five landmark pins in the patient's femur and tibia, which will act as fiducial points for registration of the preoperative plan to the actual bones.
- CT-scan the patient. Construct a 3-D bone model from CT data, with reference frames based on the landmark pins and the femoral head.
- Using a graphics workstation, plan the placement of the femoral and tibial prosthetic components.

Surgical procedures

- Immobilize the bones using specially designed fixtures.
- Use the robot to determine the coordinates of the landmark pins on the femur in order to register the femur to the preoperative plan.
- Use the robot to track the end of the femur as the femur is moved on a sphere about the femoral head. Infer the center of the femoral head for registration.
- Use the robot to guide the surgical cuts for placement of the femoral component.
- Locate the landmark pins on the tibia.
- Use the robot to guide the surgical cuts for placement of the tibial component.

Landmark pin placement

During a preoperative visit, the surgeon places four small landmark pins in the patient's bones at the knee, which serve as fiducials to correlate ("register") the CT and robot reference frames. Two pins are inserted into the distal (lower) end of the femur and two others into the proximal (top) end of the tibia. Also during the preoperative office visit, a fiberglass cast is fit around the ankle. Once hardened, the cast is carefully removed and a landmark pin is placed in the cast over the medial malleolus (inner ankle bone). These pins provide five of the six fiducials necessary to define reference frames for the two bones. The sixth will be provided by the center of the femoral head (center of the hip socket).

CT scan and preoperative planning

A computed tomography (CT) scan is obtained of the patient's leg with all landmark pins and the ankle cast (with pin) in place. Each 512 x 512 CT image, or slice, has a real width (and height) typically no greater than 200 mm, yielding a resolution under 0.4 mm per pixel. Numerous slices (typically 75) of CT data are required in order to provide a sufficiently complete model of the bones. High voxel resolution (slice spacing of 1.5 mm) is used in the vicinity of the landmark pins and the knee joint.

The CT data is read into a 486-based PC with high-resolution graphic capabilities. Edge detection algorithms are used to identify the boundaries between bone and soft tissue on each slice. Editing functions are provided to allow the user to modify the outlines. These 2-D curves are then combined into surface models representing the tibia, the femur, and parts of the pelvis and foot. A 3-D surface model of the bones is shown in Fig. 3.

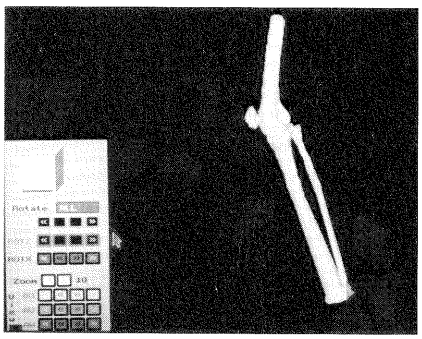

3. 3-D bone reconstruction screen.

When the 3-D model of each bone has been built, the surgeon can begin preoperative planning. Graphical software allows the surgeon full freedom to simulate almost any function that could be performed in the operating room. This includes, but is not limited to, moving the bones, rotating the joints, cutting and drill-

ing the bones, and manipulating representations of any tools necessary during surgery. By making selections from a menu, the surgeon simulates the direct placement of the prosthetic components on the bones, without introducing the approximations of a jig system. The surgeon can then view the bones from any angle and, ultimately, simulate the motion of the joint to verify the proper placement of either component. The exact criteria and algorithms for this direct process can be developed and refined, based on numerous sources including computer modeling, gait analysis, orthopaedic research, and surgical experience.

When the surgeon has determined the desired placement of each component, a command list is recorded for use by the robot controller. Included in this information are the locations and orientations in the CT reference frame of the six fiducials and of each planar cut to be made and hole to be drilled.

Femoral fiducial location

In the operating room, using a set of special fixtures, a technician immobilizes the patient's hip with respect to the robot. After making the opening incision, the surgeon places a custom-designed femoral clamp on the distal end of the femur (near the knee). The robot is then attached to the femoral clamp by means of a magnetic ball joint affixed to the robot's end-effector. (The end-effector is an integrated surgical tool, which is also equipped with drill and saw guides, a pin-finding probe, and attachment points for other accessories.) The first fiducial to be established is the location of the center of the femoral head. This point is determined by manually flexing and abducting the thigh as the robot follows and records the motion. This process, referred to as "femur tracking," estimates the center of the femoral head. Once this fiducial is established, the femoral clamp is disconnected from the robot and the femur is immobilized by a rigid fixturing arm. With the help of a small pin-finding probe attached to the end-effector, the robot is used to measure the coordinates of the two femoral landmark pins. Currently, using a calibrated robot, the accuracy of these pin-finding measurements is 0.3 mm. The robot, the surgical table, and fixturing for the pelvis, knee and ankle are shown in Fig. 4.

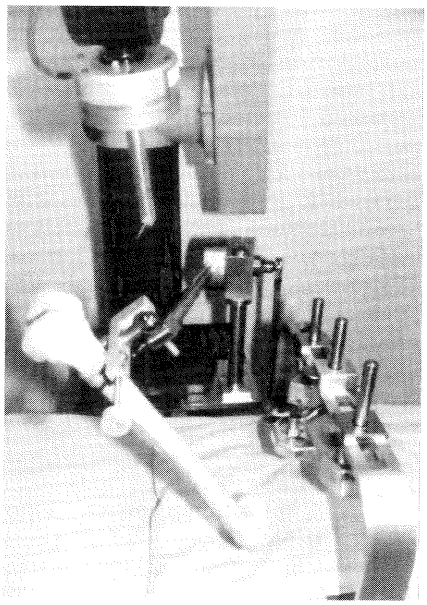

4. Robot, surgical table, and fixturing.

Femoral component placement

In order to orient the cuts for the femoral prosthetic component, a small cutting block needs to be placed on the bone. The robot is led to the approximate block location. Using the three fiducials as reference points, the robot makes a small corrective movement to position the drill guide (built into the end-effector) where the holes for the block are to be placed. After double checking the positioning, the surgeon drills the holes and then leads the robot away from the knee. The surgeon places guide pins in the holes, slides a cutting block onto the pins, and uses a powered bone saw resting on the cutting block to resect the bone.

Tibial fiducial location and component placement

A similar procedure is used for the tibia, except that the three tibial fiducial points are all landmark pins. The tibia is immobilized using the rigid fixturing arm and a rod wedged in the intramedullary canal of the bone. The coordinates of the landmark pins are located with the robot. The robot is again led to the approximate location of the holes to be drilled for a cutting block and is allowed to make a small corrective movement. After drilling the holes and making the cuts, the surgeon continues the operation, from fitting of the prosthetic components to closing the incision, in the conventional manner.

Fixturing

In order for the robot to accurately align the resections as specified by the graphics system, it is important that the patient's bones be held immobile with respect to the robot. There are several links that must be held as rigidly as possible: robot base to end of femur or tibia, surgical table to pelvis, surgical table to ankle, and robot base to surgical table. Each link has its own requirements, which will be discussed briefly.

Immobilization of the knee

The fixturing that immobilizes the distal femur and proximal tibia with respect to the robot base is the most crucial connection for the accuracy of the system. Any motion between these parts of the bones and the robot base will translate directly into errors in landmark pin localization or guide hole location. A further requirement of this connection is that it be highly adjustable, since the knee has little freedom of movement once the pelvis and ankle have been connected to the table. For this reason, a six degree-of-freedom fixturing arm with heavy duty locking joints, similar in configuration to a spherical-joint robot arm, is used (Fig. 4). The arm is designed such that no significant displacement errors result for forces applied by the surgeon during cutting operations.

Even with an extremely rigid fixturing arm, a major challenge remains to interface the arm with the bones. Only a relatively small area of bone is exposed and available for contact with an interfacing device. Further, this device must not interfere with the robot, landmark pins, cutting blocks, or bone saw. For the femur, a small bar clamp with pivoting jaws that grip the distal shaft of the bone is used. For the tibia, a rod is placed down the central canal and is firmly wedged. Each device has a protruding post that allows connection to the end of the fixturing arm.

Immobilization of the pelvis and ankle

The fixturing of the pelvis and ankle is less demanding than that of the knee, since a quarter of an inch of motion at the pelvis or ankle translates into less than a degree of rotation of the bone at the knee. However, there are no exposed bones to which clamps can be attached. Thus, sufficient clamping force must be applied exteriorly to the ankle and pelvis with special fixtures to satisfactorily immobilize the underlying bone, yet not to damage soft tissue.In immobilizing the pelvis, the fixturing must allow full rotational range of motion of the hip joint so that the leg may be flexed by the surgeon, while preventing any translational motion of the pelvis. A vacuum pack, a commercially available bag that hardens and molds to contour when air is removed, is used under the lower back of the patient during surgery to prevent pressure sore development without sacrificing rigidity. Downward pres-

sure is applied to the pelvis at three prominent areas, using anatomically contoured, foam covered aluminum blocks attached to an adjustable pressure frame ("hip-band") that connects to the surgical table (Fig. 4).

Since the distal tibia has an externally located landmark pin (in the fiberglass cast), dynamic fiducial location is not necessary, and the fixturing can fully constrain the leg. The ankle is tightly wrapped to the foot/ankle support of a modified Mark II Stulberg leg positioner that is clamped to the surgical table. The landmark pin is left exposed, and is located in the same manner as the landmark pins at the knee.

Fiducial identification and registration

Crucial to the accuracy of the system as a whole, is accurate registration of the femur and tibia to their images in the CT data upon which preoperative planning was done. Therefore, it is essential that the locations of the fiducials be accurately determined in both the CT image and on the actual bones in the operating room.

Landmark pins

The landmark pins must hold the thin cortical and the underlying soft trabecular bone of the knee joint without loosening, and also be small enough to be inserted through the skin as an office procedure. Further, the pins must be easily identified both in CT images and in the operating room by the robot. The landmark pins are modified titanium screws with a chisel point machined into the tip and a recessed cone machined into the head. A drive tool allows the surgeon to insert the pins into the patient's bone.

The five landmark pins (four in the patient's bones and one affixed to the fiberglass ankle cast) are identified in the CT data, and their locations and orientations stored in the computer. In the operating room, the robot finds each pin by having its pointer manually guided to the vicinity of the pin and then, under force control, advancing slowly along the axis of the pin until the end of the pointer is seated accurately in the bottom of the recessed cone.

Center of the femoral head

Assigning a third fiducial to the proximal femur presents a challenge. In most patients, the thigh is surrounded by sufficient soft tissue (muscle and adipose) to prevent any artificial fiducial on the skin surface from being consistently and reliably located with respect to the bone. Further, it is undesirable to subject the patient to the trauma of inserting a landmark pin anywhere but at the joint involved in the surgery. Therefore, the sixth fiducial chosen is one that is never directly touched: the center of the head of the femur. This point is found in the CT scan by identifying several points on the spherical face of the femoral head and calculating its center.

In the operating room, the robot must indirectly find the location of the center of the femoral head. With the robot clamped to the knee via the magnetic ball joint on the femoral clamp, the surgeon manually flexes and abducts the entire leg (which is able to rotate only about the femoral head) through substantial arcs, while the attached robot samples positions. The center of the femoral head can then be inferred as the center of a sphere fit to the recorded positions of the robot endpoint. It is important that the surgeon move the leg only through those motions that do not cause the center of the femoral head to move. A "map" of those motions that minimize this error has been determined through empirical studies of pelvis motion during femur tracking [2].

Robot system and calibration

The robot system is based on a 6-degree-of-freedom PUMA 560 robot with a VAL (**author: identify**) controller and a 6-axis endpoint force sensor. The surgical end-effector is mounted to the end of the force sensor. A 486-based PC communicates with the VAL controller and serves as the system's high level controller and command computer. A command script file, written prior to the surgery, defines each step of the TKR procedure to be performed in the operating room. During the operation, the surgeon communicates with the robot through pushbuttons on a hand-held control box which directs the sequencing of steps. Commands available via the command computer include large "passive" movements of the robot by the surgeon (force following), small precise adjusting movements made by the robot, and programmed sequences that find landmark pins and identify the center of the femoral head.

Critical to successful placement of the prosthetic components is the robot's own accuracy. Even with sufficient CT resolution and adequate fixturing, the promised gain of accurate component placement is not achievable with an inaccurate robot. Off-the-shelf robots are surprisingly inaccurate, and calibration is a practical way to improve their accuracy. In the rest of this section the motivation, theory, implementation, and performance of our calibration method are presented. A more detailed description may be found in [3].

Why is calibration necessary?

Since the placement of the prosthetic components is programmed off-line, the robot's absolute accuracy, in addition to its repeatability, is crucial for accurate implementation of the correct robot poses during surgery. The pose of the robot is expressed as a 6-vector (6 values) consisting of the position (3 displacements) and orientation (3 angles) of its endpoint with respect to some reference coordinate frame (for instance, the robot base). The robot's repeatability is the precision with which its endpoint achieves a particular pose under repeated commands to the same set of joint angles, and is simply a function of how well each joint returns to the required joint angle.

Absolute accuracy represents the closeness with which the robot's actual pose matches the pose predicted or commanded by its controller. Robots may have high repeatability while having low absolute accuracy. Given the joint angles, the controller of the robot computes its endpoint location and orientation. For this, the controller needs an accurate description of the robot, which involves many physical parameters such as link lengths and joint offsets. These numerical parameters make up the kinematic model of the robot. The absolute accuracy of the robot depends on the accuracy of this model.

For various reasons, the numerical values of the kinematic parameters for the robot may be inaccurate. This may be due to manufacturing tolerances, deviations such as link and joint compliance, or time-dependent effects such as gear wear and component damage. Therefore, the nominal kinematic model, which is programmed into the robot's controller, cannot accurately compute the endpoint pose from the joint angles. A practical approach to address this problem is to re-evaluate the kinematic parameters of the robot by using a calibration scheme. The implementation of such calibration schemes usually requires an instrument to measure the robot endpoint pose at various locations throughout the workspace, and a suitable algorithm for re-evaluating the kinematic parameters from the recorded pose data.

Calibration using a telescopic ball-bar

The measuring device used is a telescopic ball-bar. It is relatively inexpensive, easy to use, and highly accurate. The

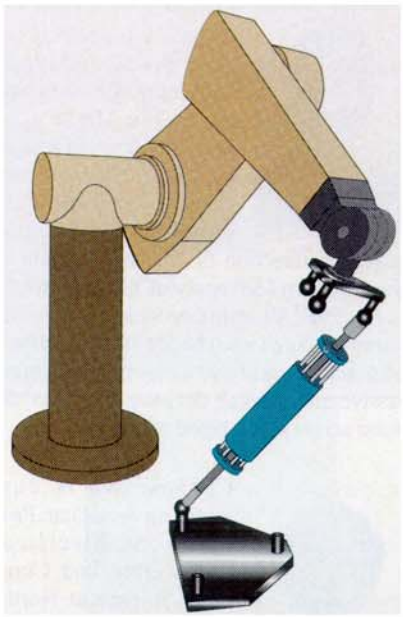

5. The calibration system with the ball-bar connected between one of three steel spheres attached to the robot endpoint and a magnetic chuck mounted on the robot table.

heart of the system is a linear voltage displacement transducer (LVDT), with a maximum travel of 7.5 cm. The LVDT precisely measures the distance of the robot endpoint from a fixed location on a triangular base plate. The system set-up is shown in Fig. 5.

The ball-bar has a magnetic chuck permanently mounted at one end, and a removable high precision steel sphere mounted at the opposite end. The removable sphere allows the insertion of extension rods, which increase the nominal length of the device in order to reach more of the robot's workspace. Additional magnetic chucks and steel spheres mate with the ends of the device to form spherical joints. In this implementation, the sphere end of the ball-bar pivots around one of three stationary magnetic chucks mounted to the robot table, while the chuck end of the ball-bar mates with one of the three steel spheres connected to the robot's moving endpoint.

The robot is programmed (using the uncalibrated (nominal) kinematic model) to move its endpoint to various poses on an imaginary shell having a radius equal to the nominal length of the ball-bar. The calibration system records robot joint positions at the various poses on this shell. However, since the parameters of the nominal robot kinematic model are imprecise, the robot endpoint will actually end up either above or below this imaginary shell. The deviation from the nominal ball-bar length is recorded by the LVDT. This deviation is known as the residual error. The aggregate sum-of-squares of the residuals (RMS error over all recorded poses) is used as an objective function, and is minimized by methodically changing the values of the robot's kinematic parameters. These new parameters are a more accurate description of the robot's kinematic structure, and increase the accuracy of the robot.

In order to identify all of the independent parameters of a robot (36 in the case of the PUMA 560), one usually obtains both position and attitude data of the robot endpoint, which requires the use of a sophisticated measuring device. The single scalar distance measured from the ball-bar falls far short of the 6-vector that would be required in order to fully determine the pose of the endpoint. This shortcoming is addressed by connecting the ball-bar between six different chuck-sphere pair combinations during the calibration. By collecting pose data on the set of six separate shells, we can identify the complete set of independent parameters.

Discussion of calibration results

The accuracy of this calibration method has been checked with two different measurement systems. First, the ball-bar was used to determine the accuracies on a larger set of shells than those used during calibration. In the second test, an interferometric laser measurement system was used to provide an independent check on accuracies through the robot's workspace.

In the first test, the success of the parameter estimation process was checked by collecting radial distance data with the ball-bar for two distinct sets of hemispherical shells. The radius of the first set was 46 cm, and the radius of the second set was 61 cm, obtained by using a 15 cm extension of the ball-bar. Each set consisted of about 800 poses.

The 46 cm radius set of shells was used to calibrate the robot. The optimal kinematic parameters were computed from the recorded poses, and reduced the RMS errors about the nominal shell radius (of 46 cm) to 0.084 mm. To check whether this new set of kinematic parameters betters the accuracy at locations not on the 46 cm shells, the RMS error on the 61 cm shells was computed, and was found to be only 0.110 mm. This result implies that the optimized kinematic model correctly reflects the robot geometry, and is not simply a best fit of the data used for calibration.

It is important to note that errors in the ball-bar lengths are only radial errors. The errors encountered during use in surgery will actually be Cartesian position errors, and are roughly four times the ball-bar error, or 0.34 mm.

The interferometric laser measurement system was used to measure errors at the robot endpoint as the endpoint moved along a straight trajectory defined by a laser beam. The measured distance was compared to that commanded by the controller. Both the uncalibrated (as delivered from the manufacturer) kinematic parameters and the calibrated kinematic parameters were used in the controller. Figure 6 shows the accumulated errors along an x-direction path in the robot's workspace. When the nominal kinematic parameters were used, the error accumulated to a maximum of 6.5 mm over the 1.25 meters traversed by the robot endpoint. When the calibrated kinematic parameters were used, the accumulated error

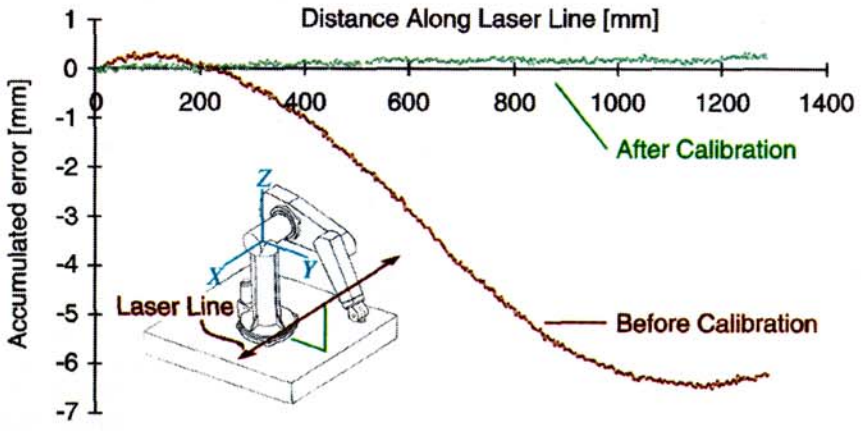

6. Accumulated position error along a laser line projected through the robot's workspace

never exceeded 0.35 mm. In fact, within the robot's workspace containing the set of calibration shells, the accumulated error never exceeded approximately 0.20 mm. Accumulated errors along paths in the y and z directions have similar values.

Future work

Preparations are being made for extensive cadaveric testing to assess the accuracy of component placement. Preliminary subsystem testing of fixturing and robot accuracy indicate that a goal of less than 1 mm of translational error and less than 1 degree of rotational error is achievable.

As the system evolves from the laboratory to the operating room, improvements to provide additional reliability, safety, accuracy, and ease of use will be incorporated. One improvement currently being developed is the use of the robot's end-effector as a cutting guide instead of as a drill guide. This approach will eliminate several manual steps in the procedure involving the cutting blocks and also alleviate inaccuracies they introduce. Another improvement is a preoperative end-effector calibration, which corrects any mounting errors between the robot and end-effector.

In the future, some elements of this system could be adapted for other surgical procedures, such as osteotomies, ligament reconstructions, and arthroplasty of joints other than the knee.

Acknowledgments

Support for this research was provided by Northwestern University, Northwestern Memorial Hospital and NSF PYI Grant DMC-8857854. Mike Brown, Dane Watkins, Rob Lentini, Boris Klovsky, and Amy Stock helped us with the fixturing prototypes. We would also like to thank James Bosnik of Drexel University for useful discussions, especially in regard to the idea of the robot endpoint fixture. Kornel Ehmann of Northwestern University drew our attention to the use of the ball-bar for calibration purposes. Kam Lau of Automated Precision Inc. was instrumental in making the ball-bar device available to us at a subsidized price.

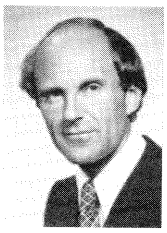

Thomas C. Kienzle III, MD, is a resident in the Department of Orthopaedic Surgery at Mt. Sinai Hospital in Cleveland, Ohio. He holds a Sc.B. in electrical engineering from Brown University, and a M.S. in electrical engineering from Syracuse University. He was employed at IBM Corporation as both a hardware and a software engineer for eight years. He received the M.D. degree from Northwestern University in 1993.

S. David Stulberg, MD, is an Associate Professor of Orthopaedic Surgery at Northwestern University School of Medicine. He is founder and director of the Joint Reconstruction and Implant Service at Northwestern Memorial Hospital.

Michael Peshkin received a Bachelor's degree from the University of Chicago, a Master's degree from Cornell University, and a Ph.D. degree (1986) from Carnegie-Mellon University, all in Physics. At Carnegie-Mellon, he worked in CMU's Robotics Institute, under the direction of Arthur Sanderson. He is presently an Associate Professor of Mechanical Engineering at Northwestern University, and serves as an Associate Editor for IEEE Transactions on Robotics and Automation, and as Editor for Conference Publications for the IEEE Control Systems Society. He is an author of 15 journal publications and one book, "Planning Robotic Manipulation Strategies."

Arthur Quaid received the Bachelor's degree in mechanical engineering in 1993 from Northwestern University. As an undergraduate, he worked on robot calibration and the development of a robot-assisted total knee replacement surgery system. He is currently working towards a Ph.D degree in robotics at Carnegie Mellon University's Robotics Institute in Pittsburgh, PA.

Jon Lea received the Bachelor's degree in mechanical engineering from Western Michigan University in 1991, and the Master's degree in mechanical engineering from Northwestern University in 1994. He is currently a doctoral student in the Department of Mechanical Engineering at Northwestern University. Previous to his graduate work, he was employed at Stryker Corporation, a maker of powered surgical tools for orthopaedic and arthroscopic procedures.

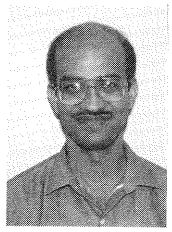

Ambarish Goswami received the Bachelor's degree from Jadavpur University, Calcutta, India, and the Master's degree from Drexel University, both in mechanical engineering. He earned the Ph.D. degree under the direction of Michael Peshkin at Northwestern University in 1993. Currently he is at INRIA Lab in Grenoble, France. His current research interests are in robot calibration, force guided assembly, programmable passive mechanical devices, computer-assisted surgery, and biped robots.

Chi-haur Wu is currently an Associate Professor of Electrical Engineering and Computer Science at Northwestern University, Evanston, IL. He received the B.S. degree in electrical engineering from National Taiwan University, Taiwan, in 1973, the M.S. degree in electrical engineering from Virginia Poly. Ins. & State University in 1977, and the Ph.D. degree in electrical engineering from Purdue University in 1980. After he graduated from Purdue, he joined Unimation Inc., Danbury, CT. During that period, his job involved in designing robot motion control algorithms and digital servo systems for PUMA robots and hydraulic-servo Unimate robots. Since September 1983, he has been with Northwestern University.

Authors' address for correspondence: Prof. Michael Peshkin, Department of Mechanical Engineering, Robert. R. McCormick School of Engineering and Applied Science, Northwestern University, 2145 Sheridan Road, B224, Evanston, IL 60208-3111. E-mail: peshkin@nwu.edu

References

1. **Paul HA, Mittlestadt B, Bargar WL, Musits B, Taylor R, et al:** A surgical robot for total hip replacement surgery. *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, 1:606-611, 1992.

2. **Lea JT, Watkins D, Mills A, Peshkin MA, Kienzle TC, Stulberg SD:** Registration and immobilization for robot-assisted orthopaedic surgery. *Proceedings of the First International Symposium on Medical Robotics and Computer Assisted Surgery*, 1:63-68, 1994.

3. **Goswami A, Quaid A, Peshkin MA:** Identifying Robot Parameters Using Partial Pose Information. *IEEE Control Systems Magazine*, 13(5):6-14, 1993.