Electric Field Pretouch: Towards Mobile Manipulation

Brian Mayton
Department of Computer Science and Engineering
University of Washington
Seattle, WA 98195-2350
Email: bmayton@cs.washington.edu

Eric Garcia
Department of Electrical Engineering
University of Washington
Seattle, WA 98195-2500
Email: eric.garcia@gmail.com

Louis LeGrand, Joshua R. Smith
Intel Research Seattle
1100 NE 45th Street
Seattle, WA 98105
Email: lllegrand@intel.com
joshua.r.smith@intel.com

***Abstract*— Pretouch sensing is longer range than contact, but shorter range than vision. Our hypothesis is that closed loop feedback based on short range but non-contact measurements can improve the reliability of manipulation. Such feedback may be particularly useful in mobile manipulation and personal robotics: mobility increases the robot's position uncertainty, which the additional feedback can overcome. In natural human environments in which Personal Robots must function, objects may be moved by people. By the time the hand of a fetching robot reaches the target, it may have moved. For human to robot hand off, the human's hand will not generally be perfectly stationary, so feedback can help aligh the robot's hand to the human's hand.**

This paper describes a series of experiments with Electric Field Pretouch that begin to explore how pretouch sensing can be used to aid manipulation.

I. Introduction

This paper presents a grasping system that is guided at short range by a sense that is not common in robotics, E-Field Pretouch. We describe two sets of experiments. The first set of experiments involves human-to-robot and robot-to-human handoff, including the use of EF Pretouch to detect whether or not a human is also touching an object that the robot is holding, which we call the "co-manipulation state." In the second set of experiments, the robot picks up standalone objects. We describe a number of techniques that servo the arm and fingers in order to both collect relevant geometrical information, and to actually perform the manipulation task (by moving the hand to the object, aligning with it, and preshaping the fingers to the object before grasping). A challenge for future work is to move from these individual examples to more general and explicit state estimation.

Our previous work on Electric Field Pretouch[11, 13] has demonstrated servoing of robot arms to align with sensed objects, and, in separate experiments, pre-shaping of fingers (on an isolated hand, not mounted to an arm) to prepare for grasping. This paper integrates these two formerly separate capabilities, and shows that they can be combined for grasping. Further, this paper explores several strategies for aligning the arm with the object (in both angle and translation) that point to a more general problem of doing explicit state estimation from pretouch measurements made in the course of manipulation.

Fig. 1. Photograph of Electric Field Pretouch manipulation system. The system includes a WAM arm and Barrett Hand. Each fingertip contains one of our custom Electric Field Sensing boards and 4 sensing electrodes. Another EF sensing board with one transmit electrode is built into the palm. The palm board also serves as a hub for aggregating sensor data from the fingertips. The palm also contains a camera, which was not used in the experiments reported in this paper. The googly eyes also were not used. A polar coordinate system used for mapping the region of where the grasping procedure succeeds is visible beneath an example object (a can).

A. Motivation

Our hypothesis is that closed loop feedback based on short range but non-contact measurements in robotic hands can improve the reliability of manipulation. Such feedback is particularly relevant for mobile manipulation and personal robotics: mobility may increase the robot's position uncertainty, which the additional feedback may overcome. Imagine a scenario in which the robot is attempting to pick up an object that is moving relative to the robot, perhaps because the robot is in motion. In the natural human environments in which Personal Robots must function, objects may be moved by people. By the time the hand of a fetching robot reaches a target location selected by vision (using techniques such as those described in [9]), it may have moved. Visual servoing[3][6] with a palm camera can allow the arm to maintain alignment with an object, but this technique will not be feasible at very close range. It is at close range that the configuration of

Robotics: Science and Systems 2009
workshop on Mobile Manipulation in Human Environments

the manipulator with respect to the object becomes a critical determinant of the success or failure of the grasp. Vision with a camera not mounted in the palm will encounter occlusion problems when the hand closely approaches the object.

Human-to-robot hand off is another scenario relevant to personal robotics in which the object may not be stationary relative to the robot hand; feedback should help align the robot's hand to a human's moving hand.

II. Related Work

Recently Hsiao et al.[2] described an optical pretouch system, with optical emitters and detectors built into the fingers of a Barrett Hand. An advantage of optical pretouch over Electric Field Pretouch is that it works with a wider range of materials. However, unlike EF Pretouch, it depends on surface color and texture and is challenged by specularity. For example, one of the failure cases described in [2], a highly specular metal object, would be ideal for Electric Field Pretouch. We believe that combining more than one pretouch modality has the potential to be very effective.

Hsiao et al. demonstrated a reactive controller for grasping using optical pretouch sensors. Hsiao's reactive controller performed wrist orientation servoing based on asymmetry of finger joint angles, a technique that we also adopt in this paper, and begin extending. In all the experiments reported here, the e-field sensor values can affect the entire arm state, not just the fingers and wrist. Many of the techniques for using pretouch sensor information presented in this paper can be applied in the context of other pretouch sensors like optical.

Another relevant point of comparison to the present paper is visual servoing.[3]. While the control and dynamics have much in common with the work presented here, a major difference is the reliance on computer vision, which provides lower update rates. Also, it is difficult to use vision at the very close ranges at which pretouch works.

Capacitive sensing has been explored in robotics in various contexts. In addition to our prior work, recently Solberg, Lynch, and MacIver presented fish-inspired underwater robot capable of localizing object using electric field sensing.[12] For several earlier instances of above-water capacitive sensing for robotics, please see [7], [8], [4] and [10]. None of these schemes were targeted specifically at sensing for manipulation.

We are not aware of prior work on co-manipulation state detection.

III. Apparatus and Methods

A. Sensing

In this subsection we describe the physics of electric field sensing, as well as the sensor hardware.

1) Physics of Electric Field Sensing: In Electric Field Sensing, an AC signal is applied to a transmit electrode (labeled TXANT in the figure). This induces an AC current in the receive electrode, labeled RCV ANT, which is amplified and processed by the Analog Front End (a current amplifier, which measures current induced at the receiver) and subsequent signal processing (in our case, an analog to digital converter and signal processing software in a microcontroller). The sensed object, drawn as a third antenna, modifies the current induced in the reader by interacting with the transmit and receive antenna.

In the case that the object is grounded, then bringing the sensed object closer to transmit-receive pair "steals" displacement current that would have otherwise reached the receiver, *decreasing* the measured sensor value. At the opposite extreme, bringing a floating object (i.e. with no coupling to ground) near the transmit-receive pair causes additional displacement current to reach the receive electrode, *increasing* the measured sensor value. The floating object is "short circuiting" the transmit electrode to the receive electrode (literally shortening the distance through the air through which the field has to propagate), while the grounded object is acting as a shunt to ground. To re-iterate, when an object is brought near a transmit-receive electrode pair, the sensor values can go either up or down, depending on the coupling of the object to ground. If the ground coupling can vary in practice, then clearly it is crucial to understand, as it can drastically affect the sensor behavior.

Note that at the frequencies at which we are operating, the human is typically well-coupled to ground, often through the shoes. (We are in the regime of AC coupling, so although there is typically no DC electrical path from your body to ground, there is usually a relatively good AC path to ground.) This in turn means that conductive objects that a person holds or touches are also relatively well-grounded.

Another important general property of the E-Field sensors is lengthscale. The range of the sensors is determined by the transmit-receive spacing. We will make use of this, performing long-range measurements to guide gross arm movements such as aligning the hand with the object, and shorter range measurements for finer finger adjuestments.

The sensors detect both conductive objects, and non-conductive objects whose dielectric constants differ from that of the air. Only the surface of a conductive object affects the sensors. For dielectrics, however, the entire bulk of the material affects the sensors. For this reason, the net dielectric constant of an object is proportional to density. Figure 2 compares the response of various dielectric objects to a conductor. Some of the dielectric objects work quite well. The ones that do not are very low density.

2) Electric Field Sensing Instrumentation: To enable the robotic hand and arm to perform electric field sensing, the three fingers of the BarrettHand were replaced with 3D-printed plastic replacements containing custom sensor boards and electrodes (both our own design). Figure 3 shows a sensor board installed on the robot. The sensor hardware is entirely contained within the plastic finger. An additional sensor board was also placed in the palm of the hand to provide another transmit channel. None of the sensing hardware described in this paper has been published before; in our prior work, we used sensor boards that were too large to be mounted in the fingers; in the prior work, only the electrodes were in the fingers. With the old hardware, it was impractical to mount

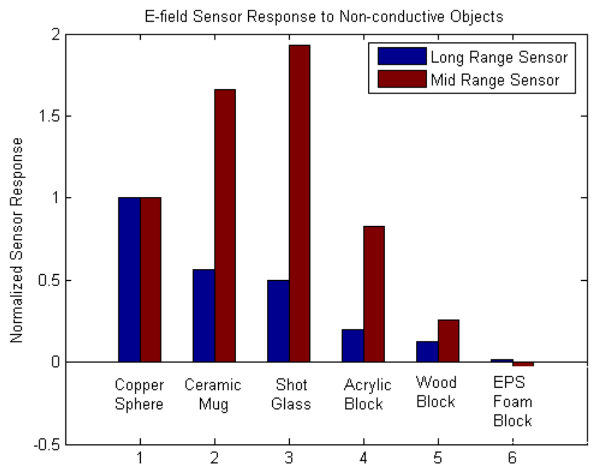

Fig. 2. Response of electric field sensors to non-conductive materials ("dielectrics"). For dielectric objects, sensor response is typically proportional to object density, which could be another useful pre-touch cue.

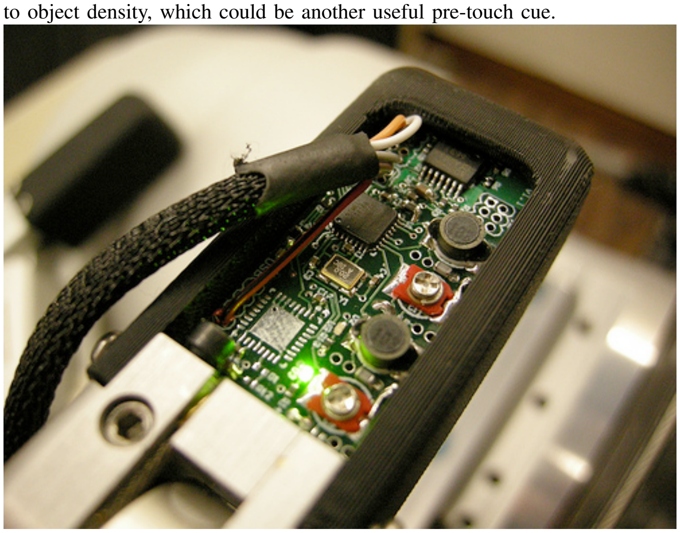

Fig. 3. Electric field sensor board installed in robotic fingertip.

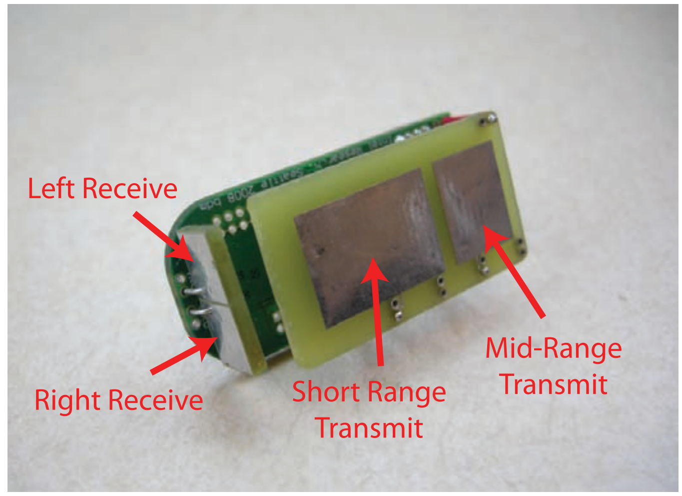

Fig. 4. Finger electrodes.

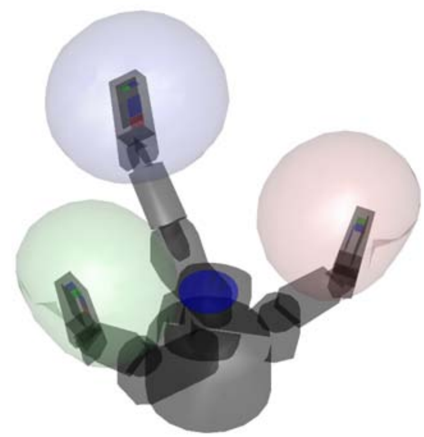

Fig. 5. Iso-signal surfaces for the hand's mid-range measurements.

the EF Sensing hand (and sensors) on the WAM arm.

3) Electrode design: The complete hand setup is capable of making 18 distinct measurements, each consisting of a transmit/receive pair. The receive electrodes are located at the tips of the fingers, and are split into left and right receivers. The placement of transmit electrode that is used determines the range of the measurement. Two transmitters are located along the inner surface of each finger. The one closest to the receivers provides a short-range measurement with high resolution but is limited to sensing about two centimeters away. The other transmitter in the finger, which is farther away from the receivers, provides mid-range measurements, with a range of about five centimeters. The transmit electrode in the palm can also be used to transmit to the fingertips, and provides a long range measurement, about 10 to 15 centimeters. (For each finger, there is a long, medium, and short range transmitter...each finger has a left and right receiver, yielding 6 measurements per finger x 3 fingers = 18 measurements total.)

Figure 4 shows the electrodes in the fingers. Figures 5 and 6 show calculated iso-signal surfaces generated by the mid-range and long-range electodes for a small test object. The iso-signal surfaces are computed by simulating the effect of a particular small test object on the sensors. An iso-signal surface is a set of locations of the test object at which the sensors return a particular single value.

B. Actuation

1) Reactive control of WAM arm: The WAM arm is controlled by a real-time Linux PC ("wambox") that provides updates at 500Hz. The sensing, inverse kinematics, and application logic execute on another PC ("wamclient") that connects to wambox by a network interface. Because of the time requirements for sensing and IK computations, wamclient provides updated commands to wambox at only around 20Hz. It is necessary to upsample from this slow, irregular set of commands to generate a set of smooth, regular commands at 500Hz.

Most of the existing WAM arm drivers upsample arm commands to 500Hz, but they require pre-planned trajectories. Given the relatively tight integration of sensing and control in our system, planned trajectories are generally not available: the next arm target location is not knowable before the next sensor value is collected. Pre-planned trajectories make smoothing

Fig. 6. Iso-signal surfaces for the hand's long-range measurements.

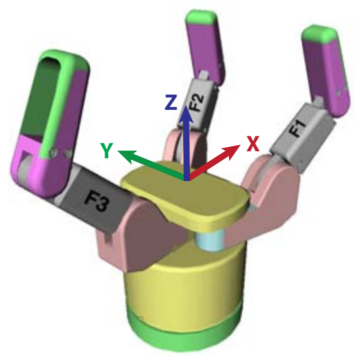

Fig. 7. The Barrett Hand, with coordinate axis.

relatively easy, since only interpolation is required; for dynamically generated trajectories, smoothing would appear to require extrapolation.

Our solution is to introduce a small amount of lag between the commands issued by wamclient and those sent by wambox to the arm. This reduces extrapolation to interpolation between previously executed commands and a future command that can be known because of the lag time. The interpolation is performed with cubic splines. The more lag allowed, the smoother the resulting trajectories. Of course the lag introduces undesirable latency, so we use a minimal lag value, on the order of 50ms. The smoothing scheme is also failsafe, meaning that if additional targets stop coming, the motion stops at the last target received.

2) Barrett Hand: The Barrett Hand has three fingers, and each finger has two links actuated by a single motor. Finger 3 is fixed to the palm, and the spread angle to fingers 1 and 2 is synchronously actuated by a fourth motor. Figure 7 shows the Barrett hand and a local Cartesian coordinate system fixed to the hand, which will be referred to as the "hand frame."

C. Control

We have explored several different control strategies for the fingers, wrist, and arm. Eventually we expect these to be subsumed by more general and principled approaches.

1) Finger preshaping using mid/short range sensors: In many grasping tasks, ensuring that the fingers contact the object simultaneously can improve the probability of successfully grasping the object. Without simultaneous contact, the first finger that contacts the object may push it out of the way, or knock it over. Simultaneous contact can be achieve by preshaping the hand to the object, i.e. commanding all of the fingers to move close to the object, without touching it. Implementing preshaping control requires a way to estimate the distance of the finger to the object. Electric field sensors in the fingers of the grasper can provide this kind of estimate

For all of the objects that we tested, the mid and short range sensors have a monotonic relationship between object distance and sensor value. Therefore, given a particular object, an e-field sensor reading set point can act as a proxy for a distance set point. In practice, it is possible to use the same sensor value set point for all objects. This is because the variation of sensor reading at a given distance for different objects is small compared to the overall range of the sensor readings. Thus the variation in distance at a given sensor reading for different objects is small is well. Further, by increasing the sensitivity of the sensor readings to distance, (e.g. short range sensors vs mid range sensors) the distance errors at a given reading are decreased.

The finger preshaping control loop used in our experiments is straight forward. The current to the finger motor (and resulting torque) is set by a PID closed loop controller that acts to reduce the magnitude of the error between the sensor reading, and the sensor set point. Each finger is controlled independently.

2) Arm servoing using long range e-field sensors: Another control mode that utilizes the long range e-field sensors to actuate the arm, but not the fingers, can be used to move the hand to a position around a stationary object, or track a moving object. This is accomplished by holding the fingers still and spread apart to give maximum position diversity. The long range sensor readings are used to create control signals in the frame of reference of the hand, and the arm is actuated via inverse kinematics. Three separate PID control loops control the motion in the three Cartesian coordinate directions in the hand frame. The controller in the y direction acts to reduce the magnitude of the difference between the sensor readings in fingers 1 and 2. When the difference is zero, the distance between the object and the sensors is equal, and the object is centered between fingers 1 and 2. Similarly, the controller in the x direction acts to reduce the magnitude of difference between the finger 3 sensor reading from and the average of the finger 1 and 2 readings. The x and y controller act independently of one another. A third controller in the z direction acts to drive the average of all three sensor readings to a predefined set point, and thus position the hand the desired distance from the object. This controller is suppressed while the x and y controllers are moving the hand. This helps assure that the hand does not bump into the object as it its moving toward it.

3) Finger encoder-based wrist rotation: The rotational asymmetry of the current finger position is computed by subtracting the encoder value for finger 1 from the encoder value for finger 2. This value is used as the input to a proportional controller that rotates the wrist in order to reduce the asymmetry of the two finger configurations. This technique was introduced in [2]. This finger encoder-based wrist servoing has the effect of orienting the hand to be parallel with the object (for various simple object shapes...for complex object shapes, the result of this rotational servoing would be harder to characerize).

4) Arm servo control using finger encoders as sensor inputs: This mode generalizes the wrist servoing described above. The arm translates the hand, as in the arm servoing section above, but the "sensor" inputs are the finger encoder values, which in turn are set by the pretouch servoing technique. Thus in this mode, the e-field sensors do not directly affect the arm state, but indirectly through the finger joint angles. [1]

In this control mode, the arm is positioned near an object to be grasped, and the finger preshape controllers (as described above) are started. As the fingers preshape, the encoder positions of each of the fingers are used as inputs to control the velocities of the arm in order to align the arm with the object and make the finger configuration more symmetric. This helps ensure that when the final gripping force is applied the object will not be displaced or rotated.

5) Finger encoder-based wrist translation: The translational asymmetry along the X-axis in the hand frame is computed by subtracting the average encoder positions of fingers 1 and 2 from the encoder position of finger 2. This value is then used as the input to a proportional controller to obtain the X component of the velocity in the hand frame.

IV. Experimental results

A. Human-robot object transfer

We have used EF Pretouch to implement human-to-robot object transfer. The human brings an object in the vicinity of the hand's long-range sensors. When the object is detected, the arm begins servoing in 3 dimensions to bring the hand into alignment with the object. (For this experiment, we arbitrarily chose an orientation for the hand. The hand maintains its fixed orientation, servos in x and y to maximize alignment with the object, and moves in and out in z to maintain a particular distance to the object.) The y error signal is the difference of the finger 1 and finger 2 long range sensor readings. The x error signal is the difference between the finger 3 sensor value, and the average of the finger 1 and finger 2 readings. When the arm is aligned with the object and the object is stationary, the system switches into grasping mode. The arm remains stationary, and the fingers pre-shape to the object. When the fingers are stationary (and in a symmetrical configuration) grasping is initiated. The grasping procedure uses the hand's EF sensors, strain gages, and encoders together to execute a reliable grasp, and detect grasp failure.

Once the robot hand has reliably grasped the object, it waits for the human to let go. This capability is what we described earlier as co-manipulation state measurement. If the human fails to release the object, the robot issues a verbal reminder: "You can let go now." Once the human lets go, the arm moves the object to another person in a pre-defined location. It prompts the person to take the object, and then waits. Co-manipulation contact detection is again used to decide when the robot hand should release the object.

The grasp control procedure is a combination of force and position control. As part of the grasping process, we want to detect contact with as much sensitivity as possible. In other words, we want to detect small amounts of force. EF Preshaping allows us to use the strain gages with more precision than would be straightforwardly possible otherwise. The strain gages in the Barret Hand fingers are affected by gravity, which can function as noise if not properly compensated, and also are subject to drift (an additional source of noise). If we were to set a contact force threshold to detect light contact of the fingers with the object, but were uncertain about the effect of gravity on the sensors, then we would have to set a contact force threshold higher. Since the fixed point of the e-field finger servoing procedure is close to the contact configuration, the effect of gravity will be similar in the two cases. Thus the strain gages can be read when the finger has preshaped to the object, but has not yet attempted to grasp it. Thus, when the fingers first make light contact with the object, this can be detected by looking for changes in the baseline strain value collected at the e-field servoing fixed point.

B. Co-manipulation contact detection

Figure 8 shows the effect of human touch on the long- and mid-range sensors for one grasp.

The long range sensors are much more susceptible to ground coupling variations than the mid-range sensors. Then a human touches the object, which increases its coupling to ground. This is possible to see in the mid-range sensor data. However, only small changes occur. For the long range sensor, human contact with the manipulated object causes drastic and difficult-to-miss changes in signal level. This makes it easy to detect human contact with an object that the robot is manipulating. We have had the robot release the object to the person when the person touches the object, and we have found the interaction to be straightforward and reliable. We made use of this technique in a high-visibility demo. In figure 10, the robot does not release the orange until the person touches the orange. The system worked so reliably that we were able to demonstrate it (arm servoing and grasping, as well as co-manipulation detection) without being embarassed by

[1]The advantage of this approach is that it removes uncertainty and complexity associated with the non-linear response of the sensors. As long as a particular setpoint (call the setpoint a null, without loss of generality) can be detected reliably by the sensor values, then the control signal (finger joint angle) needed to cause the null can be used as a sensor value, and one that may be more linear than the underlying sensing mechanism used to detect the null. This principle is used in fluxgate magnetometer sensors, whose "sensor" output is actually the control value used to null a reading on a raw (highly non-linear) magnetic field sensor.

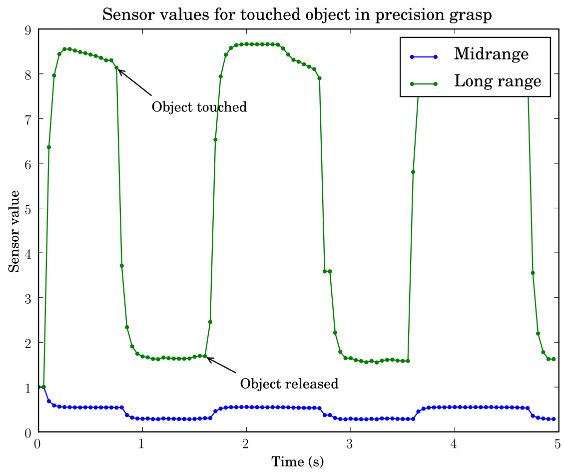

Fig. 8. Signal values from a mid-range and long-range sensor when object is held in a precision grasp and contact with a human hand is made and broken. Contact can be determined from either sensor, though the long range is much more reliable.

Fig. 9. Human handing an object (an orange) to the robot. The arm first servos to maintain alignment of the robot hand with the object or human hand. Then finger pre-shaping is used for grasping.

Fig. 10. Governor Schwarzenegger receiving orange from the robot. The robot decides when to release the object based on co-manipulation detection.

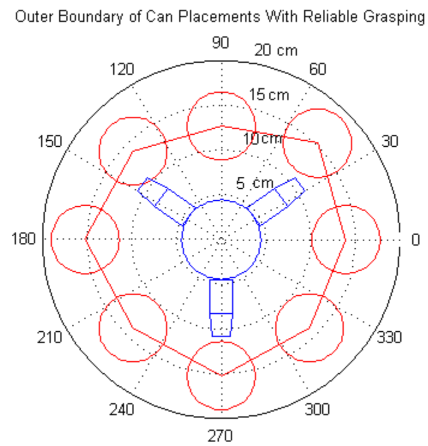

Fig. 11. “Basin of attraction” for successful object pick up. The hand always starts in the position and orientation shown, in the center. A can placed anywhere inside the red polygon will be found and picked up by the hand. A can outside the red polygon will not be detected and picked up.

demo failure in front of thousands of people, despite including untrained VIP users (Governor Arnold Schwarzenegger and Chancellor Angela Merkel) in the demo.

C. First stationary object pick up experiment

Because of the dramatic effect of ground coupling on sensor values for the long range sensors, doing straightforward arm servoing with the long range sensors would have been problematic. As the hand gets closer to the object, the floating effect becomes more pronounced. In some cases (such as a floating object at very close range), the sensor values can even reverse direction. This would lead to arm servoing in the wrong direction (away from, instead of toward, the target object).

We made use of the fact that the mid-range sensors do not enter transmit mode (“floating mode”) easily to address this problem. First, using the long range sensors with the fingers spread wide, the arm servos (in x,y, and z) until it is well aligned in x and y, and at a pre-set z distance from the object. At this point, the fingers close further, and x,y, and z servoing continues using the mid-range electrodes, which are relatively immune to impedance variations. When the hand is sufficiently well positioned, the arm stops, and finger pre-shaping begins, using the mid-range sensors. Once the fingers are pre-shaped, we initiate a grasping sequence that relies on encoders and strain gages, in addition to the E-Field sensors, to grasp the object. Figure 11 shows the “basin of attraction” within which this procedure can reliably pick up the object. This basin is roughly disc-shaped, and between 10cm and 15cm in radius. Thus if the hand can be brought to within 10cm of the object (perhaps by a longer range sensor such as a camera, or by a human command, in a tele-operation scenario), then the E-Field Pretouch grasping technique should typically be able to detect it, align with it, and pick it up.

The system has a number of parameters which must be set. The most sensitive tunable parameter is the E-Field value-set at which to switch from long-range to mid-range arm servoing. The system was tuned using a can of beans. Not only did it reliably pick up the can of beans in the basin of attraction described above, it also was able to reliably pick up an Apple, with no retuning.

D. Second stationary object pick up experiment

In this experiment, we developed a procedure that uses the encoder values of the fingers as the preshape controllers to control the arm and position the hand so that the resulting grasp is symmetric. Since only the mid- and short-range sensors are used, the ground coupling state of the object has little effect and the same procedure can be used for both electrically floating and grounded objects.

When the robot is to pick up an object, the arm is moved to a position near the object to be grasped. This position might come from a vision system such as[9], from user input via a laser pointer[5], or from a plan generated using a model of the environment, as in [1]. With a mobile robot, the actual position of the object relative to the hand might be offset by actuation and sensing errors or uncertainties in the model. In our preliminary experiments, the arm executed a preset trajectory to a specific location, and errors were simulated by moving the object.

To execute successful grasps despite these errors and uncertainties, the following procedure is used. First, the preshaping controller (as described in section III-C.1) is started using the mid-range sensors with setpoints that will keep the fingers a few centimeters from the surface of the object. A limit is also set on the maximum position to which the fingers may close to prevent them from closing too far prematurely and getting in the way as the arm is moved to be centered on the object. Once the error of the preshape controllers drops below a set threshold, the controllers for the rotation and translation of the wrist (described in Section III-C.3III-C.5) are run to optimize the symmetry of the grasp.

Once all of the controllers have stabilized and the arm has come to a stop, the preshape controllers switch to the short-range sensors with setpoints that will bring the fingers within a few millimeters of the objects surface. The controllers are allowed to stabilize again before the hand is commanded to close the remaining distance and apply gripping force. The strain gages in the fingers are used to estimate and record the encoder values at the point of contact, which can be used later to determine whether the fingers slipped.

We used this procedure to pick up a juice bottle, and a banana. The same control algorithm was used in both cases. The hand approach vector was manually provided in advance; the same approach direction would not have worked for both objects.

1) Integration into mobile manipulation platform HERB: The EF sensors were mounted on HERB, the Intel Research Pittsburgh mobile manipulation platform. Preliminary experiments allowed us to exercise the end to end system functionality. With further tuning and integration, it should be possible to use the e-field pretouch servoing on the HERB platform. HERB has cameras and laser rangefinders which can handle the long range measurements.

Note that the vision algorithm [9] currently used by HERB requires a model of any objects whose pose is to be estimated. In situations where object models are not available, getting accurate shape and pose information from vision is more difficult. In cases like these, the vision signal could do relatively simple blob tracking to crudely estimate the position of unknown objects; the e-field pretouch could take over for the final manipulation steps.

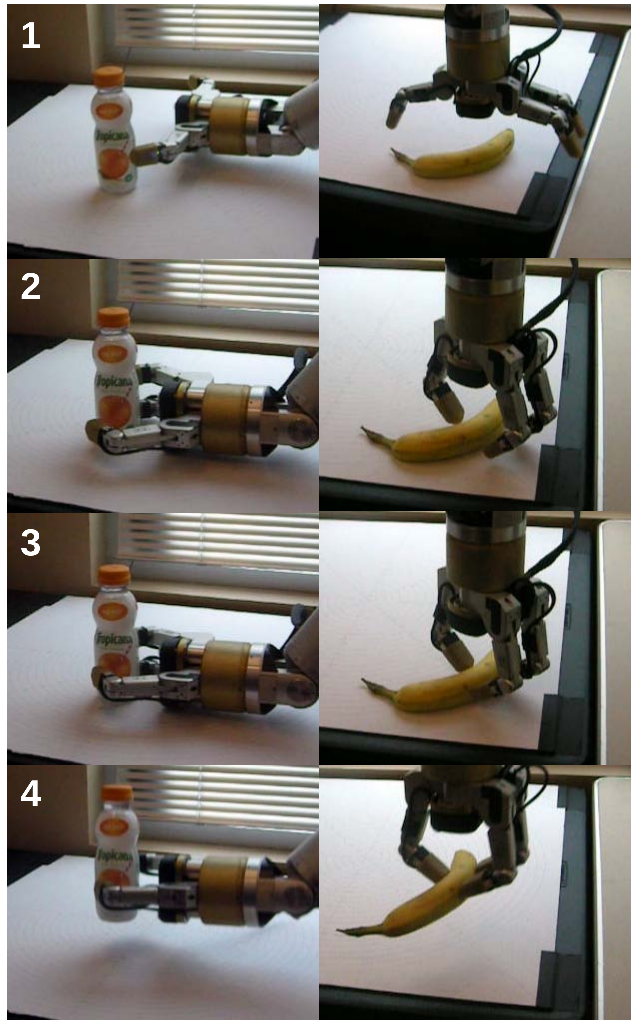

1. Hand approaches object with fingers open
2. Fingers preshape to object and wrist is translated and rotated to make grasp symmetric
3. Final preshaping is done with short-range sensors and fingers apply grasping force
4. Object is lifted from the table

Fig. 12. The EF Pretouch-based procedure picking up an orange juice bottle, and a banana. The same control algorithm is used in both cases. The difference between the two cases is the hand approach vector, which we provided manually in this case. In a full working system, the approach vector could be provided by a vision system and a planner, or by human input.

Fig. 13. Mobile Manipulation platform HERB with EF Pretouch sensing fingertips.

V. Discussion and Future Work

Our implemented system was able to pick up an object it was tuned for (a can), and also succeeded with an apple, which it was not tuned for. In a later experiment using simultaneous arm and finger servoing, it was able to pick up both a juice bottle and a banana, though different approach angles were required.

The system would certainly fail for objects that are drastically different in size from those it was tuned for. More general approaches to interpreting the sensor data are needed to allow the system to succeed with a much wider range of object geometries.

For example, currently, as the hand moves into a grasping position, we avoid needing to know the ground coupling of the object by ignoring sensor readings when they are in a regime that may be sensitive to this parameter. By combining a time history of the sensor readings and their locations with a model of the sensors, it should be possible to construct and maintain a state estimate of the object that includes its approximate dimensions, position, and ground coupling using optimal estimation techniques (e.g. Kalman filtering), which would enable more robust and general grasping capabilities.

Beyond simple estimates of object dimensions (which might be sufficient for grasping in some cases), a series of measurements could be interpreted to extract more detailed object geometry, which would allow successful grasping for an even larger class of objects. An interesting question for future work is when should primary sensor values (such as e-field sensor readings) be used, and when should secondary sensor values (such as encoder values for fingers that are servoing to null an error signal in a primary sensor).

More general approaches should also enable the system to operate successfully when multiple objects are present. The system described here would fail when presented with multiple objects.

More sophisticated state estimation approaches would also likely allow faster, as well as more general and reliable grasping. By considering many sensor measurements jointly (for example, a timeseries of sensor measurements collected as the hand approaches the object) it should be possible to produce more accurate state estimates, and thereby speed the grasping process.

Exploring the use of EF pretouch with non-conductive materials is one avenue for future exploration. Combining EF with optical pretouch[2] should allow a very wide range of materials to be sensed. The combination of sensing methods should also allow some information about material properties to be inferred.

An important step for a useful mobile manipulation system will be to combine with long range sensing modalities such as camera, laser range finder, or rfid. We demonstrated that the E-Field Pretouch can guide the hand the object from a distance of about 12cm; future useful systems will need other long-range mechanisms to get the robot hand within the 12cm basin of attraction.

Another important future step is to integrate the sensors into a full, working mobile manipulation platform. In such a setting, the benefits of EF Pretouch for overcoming additional manipulation ncertainty caused by mobility can be tested.

Videos illustratiing the systems described in this paper are available here: http://www2.seattle.intel-research.net/~jrsmith/rss09/

Acknowledgment

Thanks to Dimitri Berenson, Siddhartha Srinivasa, and Mike Vande Weghe for assistance with HERB.

References

[1] Dmitry Berenson, Siddhartha Srinivasa, David Ferguson, Alvaro Collet Romea, and James Kuffner. Manipulation planning with workspace goal regions. In *IEEE International Conference on Robotics and Automation (ICRA '09)*, May 2009.

[2] Kaijen Hsiao, Paul Nangeroni, Manfred Huber, Ashutosh Saxena, and Andrew Y Ng. Reactive grasping using optical proximity sensors. In *To appear in International Conference on Robotics and Automation (ICRA)*, Kobe, Japan, 2009.

[3] S. A. Hutchinson, G. D. Hager, and P. I. Corke. A tutorial on visual servo control. *IEEE Transactions on Robotics and Automation*, 12(5):651–670, Oct, 1996.

[4] N. Karlsson. Theory and application of a capacitive sensor for safeguarding in industry. In *Proceedings of IEEE Instr. and Msmt. Technology Conf.–IMTC 94*, Hammamatsu, Japan, 1994.

[5] C. C. Kemp, C. Anderson, H. Nguyen, A. Trevor, and Z. Xu. A point and-click interface for the real world: Laser designation of objects for mobile manipulation. In *HRI*, 2008.

[6] Danica Kragic and Henrik I Christensen. Robust visual servoing. *Intl Jour. of Robotics Research*, 22(10–11):923–940, October 2003.

[7] G.F. Mauer. An end-effecter based imaging proximity sensor. *Journal of Robotic Systems*, pages 301–316, 1989.

[8] J.L. Novak and I.T. Feddema. A capacitance-based proximity sensor for whole arm obstacle avoidance. In *Proceedings of the 1992 IEEE International Conference on Robotics and Automation*, pages 1307–1314, Nice, France, 1992.

[9] Alvaro Collet Romea, Dmitry Berenson, Siddhartha Srinivasa, and David Ferguson. Object recognition and full pose registration from a single image for robotic manipulation. In *IEEE International Conference on Robotics and Automation (ICRA '09)*, May 2009.

[10] D.J. Schmitt, J.L. Novak, G.P. Starr, and J.E. Maslakowski. Real-time seam tracking for rocket thrust chamber manufacturing. In *IEEE Robotics and Automation Proceedings*, 1994.

[11] Joshua R. Smith, Eric Garcia, Ryan Wistort, and Ganesh Krishnamoorthy. Electric field imaging pretouch for robotic graspers. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2007)*, San Diego, California, 2007.

[12] J. R. Solberg, K. M. Lynch, and M. A. MacIver. Robotic electrolocation: Active underwater object localization with electric fields. In *IEEE International Conference on Robotics and Automation*, Rome, Italy, 2007.

[13] Ryan Wistort and Joshua R. Smith. Electric field servoing for robotic manipulation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2008)*, Nice, France, 2008.