Article

The International Journal of
Robotics Research
2014, Vol. 33(5) 706–720

DOI: 10.1177/0278364913515032
ijr.sagepub.com

Mechatronic design of an integrated robotic hand

Morgan Quigley[1], Curt Salisbury[2], Andrew Y. Ng[3] and J. Kenneth Salisbury[3]

Abstract

Historically, robotic hand research has tended to focus on two areas: severely underactuated hands, and high-degree-of-freedom fully actuated hands. Comparatively little research has been done in between those spaces. Furthermore, despite the large number of robotic hand designs that have been proposed in the past few decades, very few robot hands are available for purchase on the commercial market. In this paper, we present a hand designed for minimalistic dexterous manipulation, in which every stage of the design process also considered its manufacturing cost. We discuss the various trade-offs made in the design. Finally, we present the results of experiments in which the robotic hand was affixed to a manipulator arm and teleoperated to grasp and manipulate a variety of objects.

Keywords

Multifingered hands, manipulation, grasping, mechanism design

1. Introduction

The design of robotic hands has been an active area of research for more than three decades. Indeed, over 70 hand designs can be found in the academic literature, and a similar number of designs exists in the patent literature. A summary of the large body of design research that has accumulated over the decades can be found in Bicchi (2000). Some of the most influential designs were among the earliest: the Salisbury hand (Salisbury et al., 1982) was a minimalistic non-anthropomorphic hand, whereas the Utah hand followed anthropomorphic principles (Jacobsen et al., 1986). Not long afterwards, multiple hand designs were being published each year. The number of new hand designs peaked in 2001, with nine new designs being published that year, and since then, a steady stream of new designs continues to be generated.

Not long after the early designs of Salisbury and Jacobsen, a bifurcation of design approaches seems to be found in the literature. One branch of research has pursued minimalistic, underactuated designs (Birglen et al., 2000) which seek to facilitate grasping and relatively simple manipulations: with as few as a single actuator, these hands excel at cost-effective grasping of a variety of objects. In contrast, another branch of research has pursued highly dexterous anthropomorphic designs. The objective of many of these hands is to fully replicate the capability of the human hand by providing similar degrees of freedom (DOFs), ranges of motion, and link sizes.

Although significant advances have been made in both of the aforementioned research directions, minimalistic designs for dexterous manipulation, a design space which lies between the two main bodies of prior research, has received relatively little attention. One notable exception is the Namiki hand (Higashimori et al., 2005), a relatively minimalist three-fingered hand that has been shown to be extremely agile.

Despite the large number of robot hand designs in the literature, very few robot hands exist in the commercial marketplace. This may be partially due to the many performance measures in which a robot hand must excel in order to be commercially viable, including manipulation capacity, grasp quality, sensing capability, modularity, robustness, volume, mass, speed, and cost. Many prior hand designs focused on a subset of these measures, but rarely focused on all of them throughout the design process. As a result, the transition from research prototype to commercial product may have been just as challenging as generating the original research design concept.

[1]Open Source Robotics Foundation, Mountain View, CA, USA
[2]SRI International, Menlo Park, CA, USA
[3]Computer Science Department, Stanford University, CA, USA

Corresponding author:
Morgan Quigley, Open Source Robotics Foundation, 419 N. Shoreline Blvd, Mountain View, CA 94043, USA.
Email: morgan@osrfoundation.org

Fig. 1. The hand frame and its set of identical finger modules, which attach via a mechanical fuse.

This paper presents the design of the Sandia Hand, in which we sought to balance the trade-offs between manipulation capacity, grasp quality, sensing, modularity, robustness, volume, mass, speed, and cost at every stage of the design process. Indeed, the objective was to develop a robot hand that would increase the commercially viable application space for robot hands. The design is presented in detail, including the rationale behind many key design decisions. We then present the results of experiments in which the Sandia Hand prototype performed various grasping and manipulation tasks.

2. High-level design

As mentioned previously, the design objectives of the Sandia Hand were to maximize manipulation capacity, grasp quality, sensing, modularity, speed, and robustness while minimizing volume, mass, and cost. Towards these goals, the hand was designed as a collection of identical finger modules that attach to a hand frame. Each finger module is a self-contained 3-DOF manipulator, as shown in Figure 1. The following sections will describe the high-level design decisions made in pursuit of the (often competing) design objectives.

2.1. Kinematic topology

As mentioned above, one of the design objectives of the Sandia Hand was to maximize the manipulation capacity of the hand over a wide variety of objects. In this context, we define manipulation as modifying the position, orientation, or geometry of the grasped object with respect to the frame

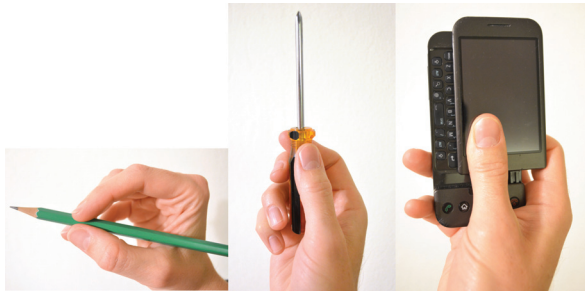

Fig. 2. Types of manipulation. Small manipulations (left), large manipulations (center), and internal articulations (right).

of the hand. This is sometimes termed *in-hand manipulation*, to draw contrast with object motions achieved by translating or rotating the wrist attachment point of a robotic arm.

For the purposes of discussing the design of the Sandia Hand, we will decompose manipulations into three different classes: small manipulations, large manipulations, and internal articulations. Small manipulations are displacements of less than 50% of the hand span in any direction or rotations of less than 120° about any axis. It is presumed that such manipulations can be realized without breaking any contact points between the hand and the manipulated object, and thus can be realized as rolling contacts between the hand and the manipulated object. A common example is the manipulation of a writing implement, as shown in Figure 2.

In contrast to the continual rolling contacts of small manipulations, we define large manipulations as displacements and rotations that require the breaking and reattachment of contact points. This assumes that the contact points do not have internally driven sliding or rolling surfaces (Bicchi and Sorrentino, 1995). A common example is the manipulation of a screwdriver through multiple turns.

Finally, internal articulations are displacements or rotations applied to the manipulated object that do not result in the uniform displacement of the manipulated object, but instead only a part of the object. Common examples include opening a mobile phone and depressing a button.

We posit that hand designs which can only realize stable grasps are a proper subset of the hand designs that can realize stable manipulations. We will begin our discussion of design considerations by determining how many contact points and associated DOFs are necessary for grasping an object. This will inform our selection of the number of fingers, finger links, and DOFs of each finger.

With the exception of platonic solids, seven contact points are required to ensure a stable form-closure grasp on any object (Markenscoff et al., 1990). Therefore, a minimum of seven links are required to realize a stable grasp on any object, which can be realized as a fixed ‘palm’ and six moving links. Although it is theoretically necessary for each link to exhibit 6-DOF positioning to grasp arbitrarily pathological objects, the resultant complexity would defeat

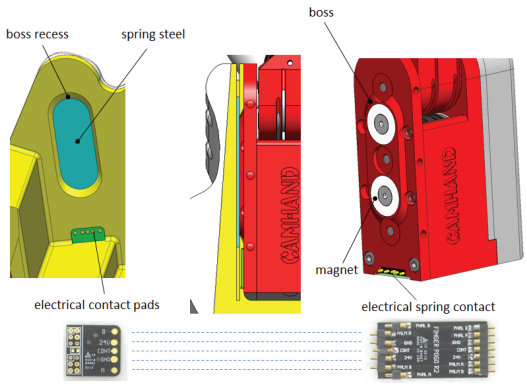

Fig. 3. Magnetic attachment of finger modules.

the goal of a low-cost, low-volume, and low-mass hand. In an effort to balance these design objectives, we made an engineering decision to attach the links to each other with 1- or 2-DOF active joints.

After surveying the full enumeration of seven-link topologies (Mason and Salisbury, 1985), we chose a more anthropomorphic topology. This decision was made to facilitate intuitive teleoperation as well as to enjoy any additional (and potentially poorly understood) benefits inherent to human hand topology. Our resulting grasping-hand design consisted of a palm with three fingers, with each finger having two links: a proximal link with two DOFs (adduction/abduction and flexion/extension), and a distal link with one DOF (flexion/extension).

Unfortunately, large manipulations and internal articulations require more links and more DOFs than those required to merely grasp an object. We surmised that an additional finger would enable the hand to realize these manipulations. The additional finger allows the hand to finger-gait a grasped object, in which one of the fingers can break contact with the grasped object while maintaining a stable grasp with the remaining three fingers. The additional finger is also capable of realizing internal articulations while the other fingers maintain a stable grasp on the object.

To simplify tele-operation, the Sandia Hand palm and its four identical fingers are arranged in a somewhat anthropomorphic topology. The palm surface normal is thus oriented perpendicular to the wrist connection axis. The thumb adduction/abduction axis directly opposes and bisects the array of three fingers, consistent with the position of the human thumb.

2.2. Robustness

The utility of a robot hand is severely limited if it is not robust to its intended environment. There are many measures that can be used to estimate the life of a robot hand, and to compare the robustness of one robot hand to another. These measures include failure loads, failure displacements, failure temperatures, failure humidity, and so on.

The Sandia Hand was designed to be robust to environmental force and displacement disturbances primarily in two ways. First, the structural components of the hand were designed to have adequate strength and toughness to withstand the loads the hand itself is capable of applying. Second, a mechanical breakaway mechanism helps protect the fingers from catastrophic failure if environmental loads exceed the capacity of the structural components. Rather than attempting to absorb such loads, in most cases the finger modules separate from the hand frame before incurring damage.

Many finger-breakaway implementations are possible. The version of the Sandia Hand described in this paper embeds permanent magnets in a chamfered boss in the front of the finger module. A strip of spring steel is embedded in the hand frame. As shown in Figure 3, the boss on the finger modules mates with a corresponding recess in the hand frame to provide 5.5 degrees of separation freedom (i.e. loads that result in the finger module being pushed into the hand frame will not result in separation). However, one challenge with this concept was fitting magnets of the adequate field density necessary to hold the finger modules in place during normal operation. Furthermore, the separation load is highly dependent upon the posture of the finger, the location of the loading point, and the direction of the loading point. When the finger is fully outstretched, and the external load is applied at the tip of the finger, normal to the long axis of the finger, the mechanical fuse will fail under smaller external loads. Furthermore, the angle of the chamfer on the boss and recess can be adjusted to tune the lateral breakaway loads.

To support finger breakaway in 5.5 DOFs, electrical connections are made with spring contact pins between the finger module and the hand frame. Spring-loaded pogo pins protrude from the front face of each finger module and press against pads on a circuit board attached to the back of the palm to conduct power and data between the palm and each finger module. As the finger disconnects during a mechanical overload event, these pogo pins simply slide or otherwise unload from their respective contact pads. The circuitry immediately downstream and upstream of this spring-loaded connection is designed to handle the transient voltages associated with momentary connections or shorts of these pins to their neighbors. As a result, mechanical breakaway can occur without mechanical or electrical damage to either the finger or palm subsystems.

An important feature of this architecture is that failure of a single finger will not disable the entire hand. The damaged finger can be quickly removed and replaced, or the hand can continue to operate without the damaged finger(s).

2.3. Modularity

As previously discussed, the strictly modular design of the hand offers robustness benefits. However, there are other

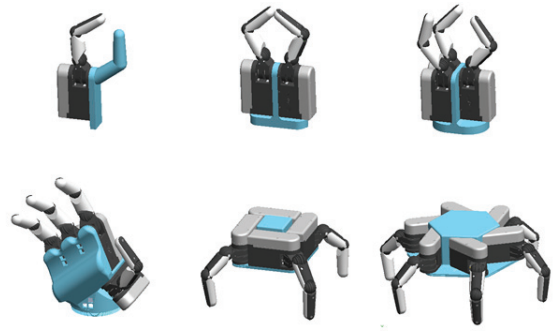

Fig. 4. Hand frame variations.

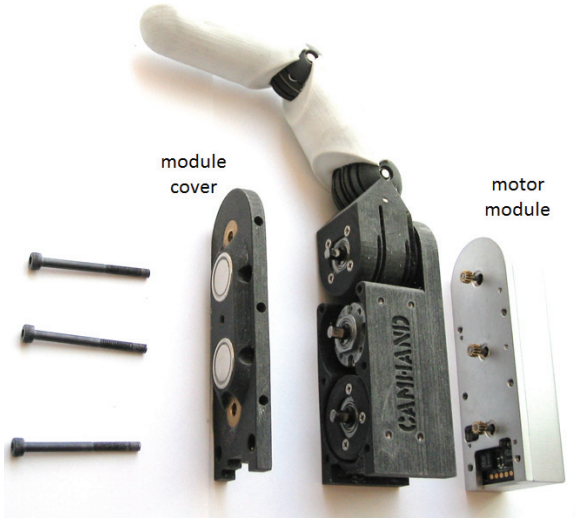

Fig. 5. Internal modularity of a finger module.

important benefits as well. For instance, each finger module socket on the hand frame could accept other devices with the same interface. These devices might include power tools (e.g. screwdrivers, drills, etc.) and sensors (e.g. particle counters, cameras, etc.). Modularity can thus enable some tasks to be easily completed by attaching a specialized tool, rather than requiring the hand to grasp and operate a tool designed for operation by the human hand.

Modularity also enables users to experiment with alternative kinematic topologies with relative ease. Only the hand frame, a relatively simple component with no moving parts, would need to be redesigned in order to realize a hand with an entirely different kinematic topology. Figure 4 shows a few example topologies that one could realize with unchanged finger modules.

Finally, on a more practical note, the modular design facilitates the replacement or upgrade of finger modules without requiring the replacement of the entire hand, or even its removal from the robotic arm holding the hand.

In addition to the hand-level modularity of removable fingers, the individual finger modules were also designed to exhibit internal modularity. A finger module consists of a cover, a tendon drive system, and a motor module, as shown in Figure 5. This design allows a damaged motor module to be replaced without rebuilding the entire finger module, and also facilitates independent upgrades or customization of either the tendon drive system or the motor module. In addition, internal modularity permits customization of the cover design for integration and attachment to a variety of hand frame interfaces.

3. Low-level design

This section describes the implementation of the major subsystems of the hand which realize the high-level design presented in the previous section.

3.1. Hand frame

The hand frame has no moving parts. Instead, its primary purpose is to multiplex power and data to the fingers and its own sensors. Simple state machines govern the voltage of each finger socket: until a finger is operating normally, the socket is maintained at a low voltage to reduce the risks associated with exposed conductors. Once communications are established with a finger, the hand frame increases the voltage and current limit of the corresponding finger socket to provide sufficient power for the finger motors.

To manage the considerable data traffic of the hand, a field-programmable gate array (FPGA) system was implemented. The hand frame provides a gigabit Ethernet data link, multiplexing the various internal sensor streams and routing traffic across its four internal finger data busses.

Finally, the hand frame has an internal sensor suite including a binocular camera array, an inertial measurement unit (IMU), and a large tactile array similar to those on the finger modules, which will be described in detail in Section 3.6.2.

3.2. Actuation

Minimizing the volume of a robotic hand is tremendously difficult, particularly if the design parameters require that all actuators reside in the hand itself. Human hands are slender in large part because many actuators (i.e. muscles) reside in the forearm. Indeed, if the human hand were to contain all of its actuating muscles, it would be substantially larger and less capable of navigating in tight spaces. The choice between using internal or external actuation is thus perhaps the largest single driver of the volume of a robotic hand. Despite the added design challenge, our tradeoff analysis strongly valued the ability to mount the hand to unmodified, pre-existing manipulator arms, and thus we chose to use internal actuation.

Numerous design decisions impact the volumetric size of an actuation system, the first of which is motor selection. Although inrunner brushless motors have higher power densities than outrunner brushless motors, inrunner motors typically require another stage of mechanical speed reduction to deliver the same form of output power (force–velocity

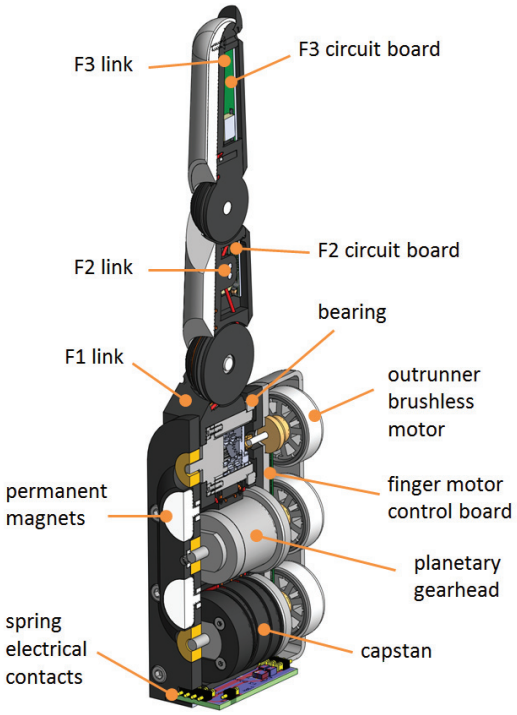

Fig. 6. Actuation and mechanical power conversion approach. Ring gears of the lower two planetary transmissions rotate and serve as tendon capstans.

ratio) to the fingers. To minimize the total actuator volume, we selected an outrunner motor, in particular, the Nidec 20N Type 2 (Nidec America, St. Louis, MO), which has a power density of 1.18 W/mm^3.

The second design decision affecting actuator volume is the mechanical transmission. If the transmission is inefficient, it will require a larger motor to provide the same output torque. Furthermore, some mechanical transmissions are more volumetrically dense that others. Worm and standard spur gears, for instance, result in a far larger transmission volume than a planetary gear transmission. Because planetary gear transmissions have high torque and reduction densities, and are readily available at mass-market prices, we selected a planetary gearhead for our reduction mechanism. In particular, we selected the Hennkwell PK22G2150-231B gearhead (Sihu Town, Changhua County, Taiwan).

To further reduce actuation volume, the ring gear of the planetary gearhead was used as the drive capstan for the finger tendons. Typically, the output shaft of a transmission is used to drive a capstan to spool the tendons, but this topology causes the capstan to drive up the overall actuator volume. If the ring gear *is* the capstan, volume can be significantly reduced. To realize this concept, we replaced the rigid attachment between the ring gear and the motor with a bearing interface and locked the output shaft relative to the finger module, as shown in Figure 6. Thus, when the motor is driven, the ring gear provides the power output.

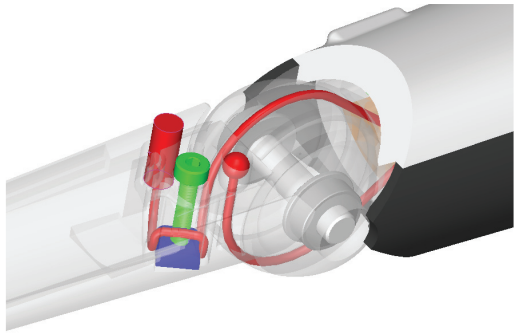

Fig. 7. Tensioning mechanism.

In order to maximize the flexibility of finger module arrangements, we arrayed the three actuators in a vertical line. The resulting actuation system can deliver a continuous force of ±10 N at the tip of a fully outstretched finger in flexion/extension, and ±4 N in abduction/adduction. No-load maximum joint speeds are ±4.5 rad/s.

3.3. Tendons

The desire for slender finger volumes required that the actuators be mounted proximal to the fingers. As such, translational mechanical power transmissions were required to transmit torque from the actuators to the finger joints. To reduce complexity, the adduction/abduction joint was designed co-axially with its actuator, allowing direct attachment to the ring gear.

However, the two distal joints required tendons. The middle actuator in the base of the finger module drives the proximal flexion/extension link, and the bottom-most actuator drives the flexion/extension of the distal-most link. The tendons were implemented with a stainless steel wire rope assembly connecting the links to their respective actuators. The specific routing path of the tendons can be seen in Figure 17. This tendon routing results in mechanical power coupling between links. Careful design of the coupling behavior enabled identically sized actuators to be well matched for fingertip output loads.

The method for tensioning the tendons is shown in Figure 7. One end of the tendon has a swaged ball fitting, while the other end has a copper crimp fitting to facilitate assembly. The tendon routes up, over, and back down a captured square nut. Turning the screw pulls up the square nut and takes up tendon slack. Note that this mechanism reduces slack by twice the travel distance of the square nut.

The resulting ranges of motion are ±90° for all joints, with an extended range of +115° for distal flexion.

3.4. Thermal management

Each motor, along with its drive circuitry, generates about 5 W of thermal power when applying maximum continuous grasp loads. To sink this heat away from the motors,

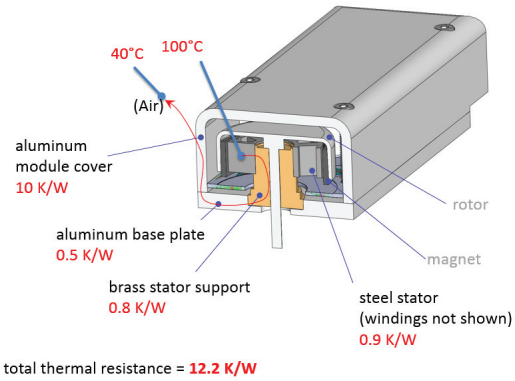

Fig. 8. Thermal aspects of motor module design.

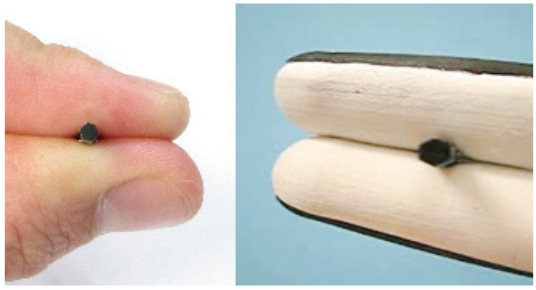

Fig. 9. Soft tactile pads allow conformal grasping of small objects.

the stator of each motor is thermally connected to the motor module cover. The thermal path includes the stator, a brass stator support, the aluminum base plate of the motor module, and finally the aluminum motor module cover, from which heat is passively removed via convection. This configuration provides 5 W of thermal dissipation for each motor when the rotor is at 100°C and the ambient temperature is 40°C, and is illustrated in Figure 8.

3.5. Contact geometry

We have observed that fingers with round cross-sections are preferred as they maintain a stable grasp while performing manipulations: round cross-sections can reduce finger-to-object displacements while rolling across the finger surface. However, a disadvantage of this approach is that the contact area can be smaller over the entire manipulation phase. In order to address this, we integrated soft coverings, or skins, on the contact surfaces (Barbagli et al., 2005). As demonstrated throughout the animal kingdom, biological manipulator surfaces such as hands, feet, paws, tentacles, etc., conform to the object or environment. Flexible surfaces create far larger contact patches than those created by rigid manipulator surfaces, allowing for partial or full envelopment of small objects, as shown in Figure 9. This effect is particularly pronounced when small objects must be manipulated in fingertip grasps, such as writing implements or keys. Even when handling larger objects, however, grasp stability is greatly improved when the contact patch is a large surface rather than the line contacts or even point contacts produced when rigid manipulator surfaces attempt to grasp hard objects.

From an engineering perspective, however, it is often difficult to fabricate a surface which exhibits conformal properties while being mechanically robust. To address this challenge, a multi-layer silicone finger skin was developed. Conformal properties are provided by a relatively thick layer of very soft Shore OO 10 silicone gel (Ecoflex 10, Smooth-On, Inc., Easton, PA). This soft gel is coated by a thin and somewhat stiffer coating of Shore A20 silicone (DragonSkin 20, Smooth-On, Inc., Easton, PA). Because these two materials are of the same type of cure, the outer skin bonds readily to the inner skin. The result is somewhat analogous to animals, where relatively stiff layers of dermis and epidermis protect the soft layers of flesh underneath. In both cases, the skin is highly compliant in the direction of its surface normal, but significantly stiffer under shear loading conditions. This allows shear forces to be exerted on objects and the environment, even while the manipulator surface is conforming to the object or environment. This construction was also used for the palm skin.

3.6. Sensor suite

The hand is equipped with a variety of proprioceptive and exteroceptive sensors. The proprioception capabilities include joint-angle, tactile, strain, and thermal sensing schemes, whereas the exteroceptive sensing capability is provided by a small camera array. These sensing modalities will be described in detail in the following sections.

3.6.1. Joint encoding. The structural and corresponding volumetric constraints of robotic fingers make it difficult to sense their joint angles. The meshing of the finger knuckles causes each link to have less material at that interface, resulting in a weak point. Reducing the material further to accommodate conventional joint encoders would compromise the robustness of the finger. Instead, a combination of MEMS inertial sensors and forward-propagation of motor encoders were employed to create a compact, low-cost joint-sensing solution. 3D accelerometers were placed on the circuit boards in the base of the finger module and the distal and proximal phalanges, as shown in Figure 6. In static conditions, these accelerometers measure the direction of the gravity vector in each respective inertial frame.

By projecting the difference between the direction of the gravity vectors onto the constraints of the kinematic chain, the joint angles can be inferred in many cases. The F2–F3 joint angle is relatively easy to infer, as accelerometers are positioned on either side of this 1D joint. As such, the joint is only difficult to measure when the joint axis approaches the vertical. In contrast, the two proximal joints (F0–F1 and F1–F2) are significantly more difficult to estimate. Because

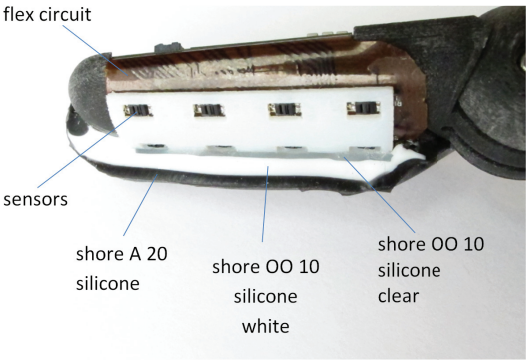

Fig. 10. Section view of the finger skin showing optical tactile sensors.

of the kinematic advantages of reducing the offset between the two proximal orthogonal joints, F1 was made as compact as possible. Its interior volume is fully occupied by four internal cable pulleys and the shaft and bearings for the F1–F2 joint. As a result, it was prohibitively difficult to place an accelerometer in this link, and thus the joint angles for the two proximal joints must be inferred by only observing the gravity vector in F0 and F2. The closed-form solution for the two proximal joint angles as a function of the accelerometer readings in the F0 and F2 links is similar to the solution popularized for inverse kinematics of similarly offset robot configurations. It results in at most four solutions, and due to joint limits, typically only one or two of these solutions is kinematically possible. We note that this sensing method is not applicable in orientations in which finger joint axes are approaching the vertical, and that it requires a stationary reference frame.

Once the initial finger posture is estimated, hall-effect sensors used for brushless motor commutation can be used to estimate the joint angles by forward-propagating the rotor positions through the various transmission elements. The theoretical positioning resolution of such hall-state encoding, however, is naturally degraded by the backlash inherent in low-cost planetary gearheads. The resulting fingertip positioning uncertainty is approximately 1 cm when fully outstretched.

3.6.2. Tactile sensing. Robotic manipulation, whether for simple grasping or complex in-hand manipulation, involves managing fingertip forces while maintaining contact with objects. As such, high-resolution tactile data can be extremely useful. Towards this objective, and with the overall design goals of low cost and robustness in mind, a novel tactile scheme was developed and implemented.

As described in the previous section, the finger pads are constructed using a multi-layer principle to exhibit mechanical stiffness against shear loads and mechanical compliance under normal loads. By measuring the deflection of the soft inner layer, the normal contact forces can be estimated.

To observe this deflection, we implemented a concept from our prior work (Lincoln et al., 2012). The Shore OO 10 silicone was split to first contain a 1 mm thick clear layer, after which the balance was pigmented white. Consistent with our prior design, an array of transflective photosensors were embedded below this clear layer, as shown in Figure 10. Transflective photosensors are comprised of an LED–phototransistor pair inside a single package, with a vergence angle such that the photocurrent varies with both reflectivity and proximity of reflective objects within a few millimeters of the device. The Shore OO 10 silicone above the clear layer was colored with a white pigment to create a reflective surface that moves closer to the sensors with normal surface loads. The variance in proximity of the white layer produces a varying photocurrent, which is passed through a transimpedance amplifier and low-pass filter before being digitized by a 16-bit analog-to-digital converter.

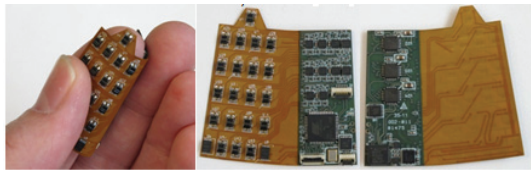

Fig. 11. Rigid-flex circuit assembly to provide tactile sensing.

Rigid-flex circuit boards were created to fit this circuitry into the space constraints of the robotic fingers. Rigid-flex constructions allow for high-density, multi-layer circuitry on a portion of the assembly, with a subset of the copper layers then continuing outside the fiberglass core and being covered by flexible polymide film. As shown in Figure 11, the rigid portion includes the vast majority of the components and is routed on six layers, whereas the two-layer flexible portion includes only the photosensor array. When installed into the robotic finger, the rigid portion resides inside the finger volume, and the flex portion wraps around the slightly conical outside of the finger core, which is covered with a protective 'window frame', to secure the photosensors against shear loads, and then is covered with the multi-layer silicone skin. This sensor concept was also implemented on the palm skin in the same fashion.

By varying the durometers and thicknesses of each respective silicone layer, a variety of sensor characteristics can be tuned, such as sensitivity, range, and mechanical toughness. For the robotic hand described in this paper, the layer thicknesses and durometers were chosen experimentally to seek a balance between these properties to allow sensing of hand-held tool manipulation. A representative plot of the raw sensor response to repeated cycles of loading and unloading a 2 g US penny is shown in Figure 12, demonstrating that these 2 g loads are far above the noise floor of the sensor.

3.6.3. Coarse contact sensing. While tactile sensors exist on all of the surfaces which are critical to grasping and manipulation, a means of determining forces elsewhere on the fingers can be useful, for example to estimate forces on

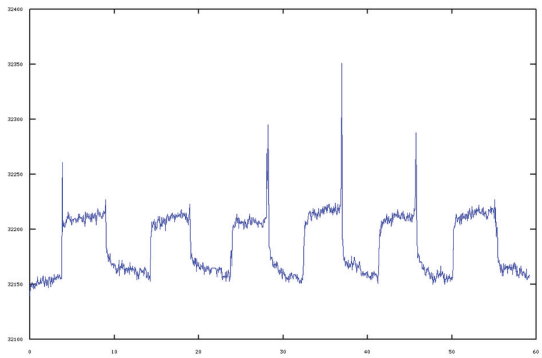

Fig. 12. Raw sensor response of repeatedly loading and unloading a 2 g US penny onto the skin assembly.

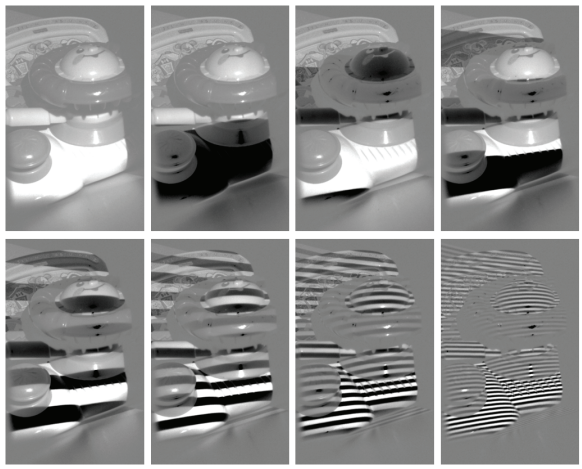

Fig. 13. Difference images produced by a laser pico projector affixed to the robotic hand.

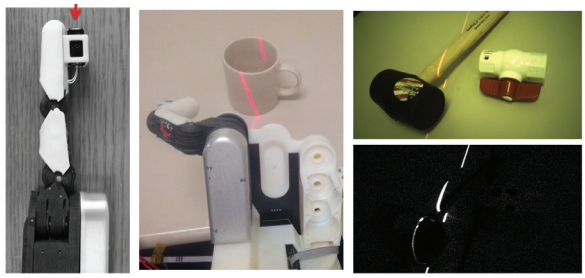

Fig. 14. LLG mounted to a robotic finger.

the back of the finger. To sense these interactions, a single strain gage was affixed to the center of the F2 link. This gage responds to bending loads about J2 and tensile/compressive loads along the long axis of F2 applied at any point distal to the gage. Though it does not facilitate independent estimation of these loads, the sensor is primarily intended to provide the simple information that contact is happening in locations other than on the skin. Because of space constraints in the F2 link, the strain is measured using a single gage in a quarter-bridge configuration.

3.6.4. Visual sensing. Although tactile and force sensing can guide motions and forces after contact, a key challenge in robotic manipulation is predicting *where* contact forces should be applied before contact is made. Decades of research and a wide variety of approaches have been applied to this problem, with many approaches utilizing vision-based data. The conventional implementation of a vision sensor, however, is near the base of the manipulator arm (e.g. the robot head). This placement is problematic for two reasons: position estimation errors and occlusion.

If visual or depth data is obtained from head-mounted sensors on a humanoid robot, this information must be propagated through typically nine joints (pan/tilt head, followed by a 7-DOF arm) to provide data in the frame of the end effector. Angular errors and coordinate-frame misalignments are cumulative, resulting in errors on the order of 1 cm even on precisely machined, carefully calibrated humanoid robots. Furthermore, a head-mounted sensor will suffer from occlusions created by the robot's own hardware. As a result, grasps of objects typically involve placing at least one contact point on a region of the object that cannot be directly perceived by the head-mounted sensor from a single viewpoint, which can be a source of errors.

Placing visual sensors in the hand simplifies these challenges substantially. Ever-smaller sensor systems and support electronics have made in-hand vision systems easier to integrate. Further improvements in contact-point planning can be obtained if depth data is generated from the vision system. A variety of vision-based depth-sensing modalities have been explored in the literature. Passive stereo systems are notorious for failing in sparse artificial scenes having low texture, and following the reasoning of Lim (2009) and Konolige (2010), we explored the integration of active-sensing modalities. Specifically, we explored the use of a laser pico projector, laser line generator (LLG), and laser speckle generator (LSG) as methods to inject light and generate depth data, comparing those data to passive stereo depth data.

The pico projector was used to create difference images of horizontal bars at various scales, as shown in Figure 13. Pixel-wise depth estimates are then obtained by classifying each pixel of each frame as {0, 1, indeterminate}, and converting the resulting binary string to its unique plane emerging from the projector. The intrinsic calibration of the camera then produces a ray, which intersects this plane and produces a 3D estimate. While this method preserves sharp depth discontinuities due to independent pixel-wise estimation, eye-safe standards restrict the brightness of this specific projector technology to be too low for normal ambient lighting conditions.

Significant gains in signal-to-noise ratio can be obtained by using an LLG, since the eye-safe optical energy of the laser is collapsed onto a 1D line, instead of a 2D surface. The LLG was mounted on the back of a finger (Figure 14) which was then moved to sweep the laser across the scene,

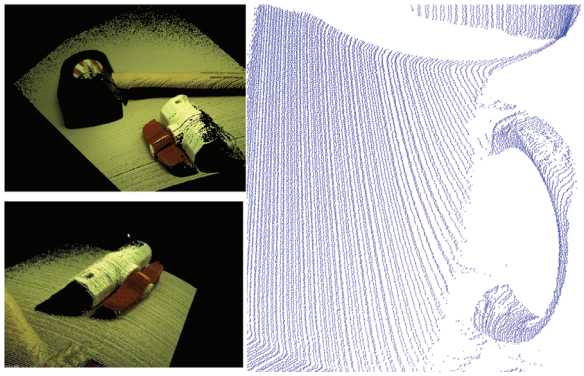

Fig. 15. Fine detail of a finger-mounted laser line scanner.

as in Butterfass et al. (1998). Camera images were then processed following established image-processing techniques for laser line scanners. While the use of an LLG enables perception with fine detail, scans are time-consuming. The scans shown in Figure 15 were acquired in approximately one minute and assembled from 300 positions of the finger-mounted LLG. An LSG, which produces unstructured light, was found to produce depth data much faster than the LLG, with a significant reduction in data density and accuracy.

A comparison of these three methods is shown in Figure 16 on a scene of a potential future application of the robotic hand: grasping a coffee mug from a table as part of a 'tidying up' task. As is common for indoor workspaces, the scene offers little texture for passive stereo block-matching, hence producing few points on the surface of the coffee mug itself. Laser line scanning (Figure 16, right-most column) produces hundreds of thousands of points, but at the cost of acquiring 300 images. Injecting texture into the scene via laser diffraction (Figure 16, middle column) offers a useful intermediate level of point cloud density with single-frame point-cloud acquisition. Based on these qualitative results, we found the LSG to be preferable for our application. The LSG was mounted on the hand frame, just below the palm. This places the LSG on the perpendicular bisector of the stereo baseline.

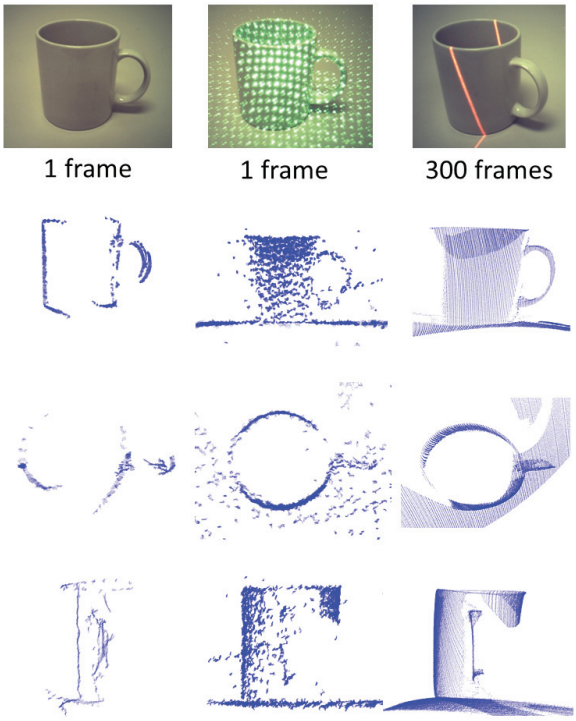

Fig. 16. Comparison of three depth-mapping techniques. Shown are raw images and depth data for passive stereo (left), laser-speckle-generated unstructured light (center), and laser-line-generated structured light (right).

3.7. Minimizing the electrical cable assemblies

Electronic cable assemblies are a common failure point in robotic systems, due to the movement of links with respect to one another. Such movement flexes electrical cables and can cause cables to be snagged by the other links or the environment. To minimize flex failures, service loops with large bend radii are usually planned at the design stage. To minimize snags, these service loops are secured close to the links of the robot. These design practices are not well suited for compact high-DOF robotic hands, as the service loops often require significant volume. To address the mechanical failure modes of electrical cables in the robot hand, we invested significant design effort in the elimination of electrical cable assemblies from the design.

We found that the mechanical cable assemblies that transmit torques to the finger joints could also be used to transmit DC power and multi-megabit half-duplex data between the finger module base and the phalanges. The advantage of this approach is that a separate set of flexing and possibly exposed electrical conductors is eliminated. While the cables were necessarily made of steel in order to perform their mechanical function, their electrical properties (resistance) were adequate and the fatigue properties were far better than those of copper wire assemblies. Unlike copper, steel will never fail in fatigue if the loads are maintained below a particular level (fatigue strength).

The complete conductive path starts at the finger motor control board and ends at one of the two phalanges. A small circuit board with electrical spring contacts creates continuity between the finger motor control board and a set of copper wires embedded in the finger module cover. Those wires are pressed in between the plastic of the module cover and a brass insert placed into the cover. Continuity is achieved between the embedded copper wire and the output shaft of the gearhead through the brass insert via a set screw. The conductive path then leads from the output shaft, through a bearing, and onto the ring gear. Continuity between the ring gear and the mechanical cable was achieved by removing insulation from a small length of the mechanical cable and sending a flathead screw through the middle of the cable braid and into an internal thread on the ring gear. Finally,

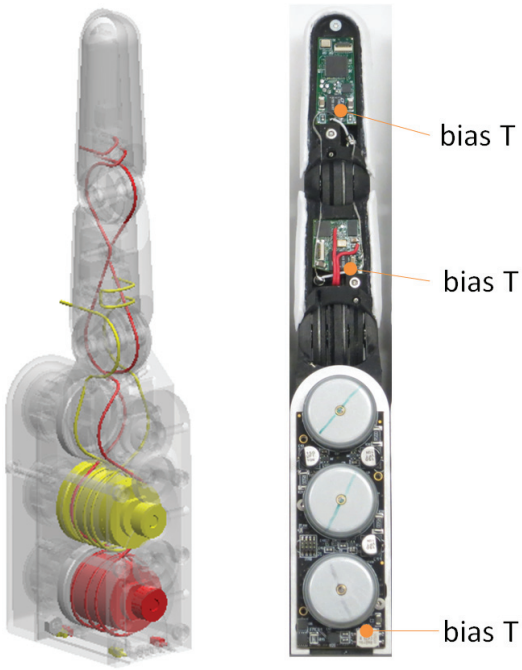

Fig. 17. The conductive paths between the phalanges and the finger motor control board.

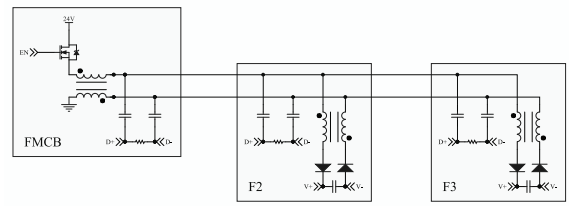

Fig. 18. Simplified schematic of the multiplexing of power and half-duplex data over the pair of conductors running the length of the finger. RS-485 transceivers are connected to the D+/D− nodes; F2 and F3 power supplies are connected to the V+/V− nodes. Bus power is supplied from the finger motor controller board (FMCB) (left).

continuity between the mechanical cable and the phalange board was accomplished by using a copper crimp on the mechanical cable and soldering one end of a copper wire to the crimp and the other end to the phalange board.

Each finger module has two mechanical cable assemblies. One assembly carried power and non-inverted data, while the other served as the ground return and carried inverted data. To bring power to the proximal phalange, and ground to the distal phalange, we again used steel cable because of its strength and flexibility in very thin gage wire, and crimped copper connector fittings to the end of the wire. The connector fittings were inserted into the corresponding connector housing, while the wire was wound around and inside the knuckle and passed to the other phalange. The electrical resistance of the conductor chain totaled 1.5 Ω. This would potentially be an issue if this transmission line were intended to drive motors. However, since only sensors and relatively energy-efficient microcontrollers are receiving power, the losses are manageable.

The potential for intermittent connectivity is addressed by capacitors on the distal sensors, allowing them to maintain power through the momentary glitches. Data dropouts are handled by a protocol in which each packet is protected by checksums and slightly spaced in time, to allow for fast and unambiguous re-synchronization of all transceivers.

The multiplexing scheme employed to transmit both power and half-duplex data across the pair of conductors consists of two sections. First, a differential-mode choke is used on each board to attenuate the high-speed differential-mode signals and produce a clean power line from the conductor pair. Second, RS-485 transceivers are capacitively coupled to the conductor pair, to readily pass the high-speed signals to and from the conductors. Each transceiver has a transmitter-enable line that is controlled by its corresponding microcontroller. Similar to USB and other half-duplex architectures, bus traffic follows a simple master/slave scheduling method to prevent collisions. A simplified schematic diagram is shown in Figure 18.

To prevent either inductor saturation or overloading the active transceiver's output stage, the microcontrollers in each module were specifically selected to include Manchester encoders. The Manchester encoding scheme is a simple method of removing DC bias in a data stream by replacing each '0' or '1' bit with a '0-1' or '1-0' pair of chips. Because the chipping rate is now twice the data rate, effectively half the bandwidth has been lost. While this condition is considerably improved in more contemporary 8b/10b coding standards, space constraints in the phalanges, the desire to use commodity low-cost microcontrollers, and the sufficiency of its throughput for tactile data transmission resulted in its selection.

In addition to the wiring reduction for communicating with the phalanges, we found that the number of loose conductors used to interface with the motors could be greatly reduced, significantly reducing the volume of cable and connectors inside the finger modules. Rather than interfacing with the motors via the manufacturer's connector, we interfaced at the motor winding leads. In order to accomplish this simplification, we integrated the hall-effect sensors directly into our driver board. This enabled us to reduce the interface conductor count from eight to three and have much shorter conductor lengths. The difference in volume required for the motor system can be seen in Figure 19.

The design challenges and approaches taken to address the motor and phalange wiring issues were somewhat specific to the design environment of a robotic hand with many DOFs. There are other interconnection challenges in the hand, however, that are common to any high-density electronic system: connecting the various planes of circuitry in the palm. These circuit boards do not move with respect

Fig. 19. Comparison of motor wire volumes between design prototypes.

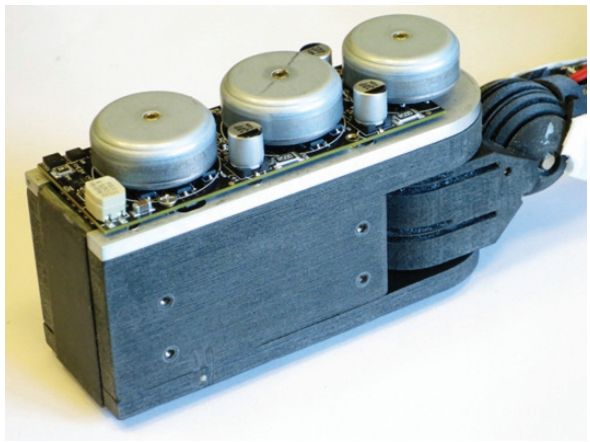

Fig. 21. Stackup of outrunner brushless motors, controller board, and heatsink.

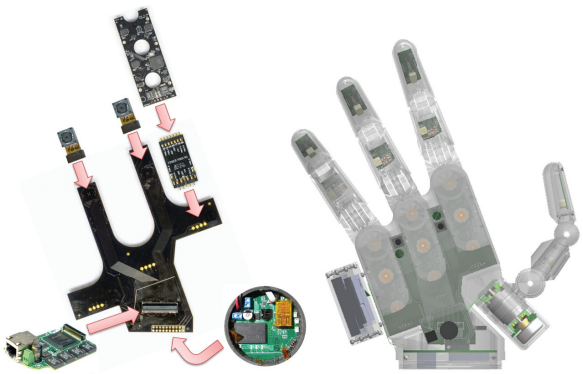

Fig. 20. Electrical and mechanical co-design allows circuit boards to blind-mate during installation, eliminating internal wiring.

to each other, and hence the standard techniques of high-density dual-row connectors can be successfully employed. The cameras, palm board, FPGA board, and power board all interconnect using 0.5 mm pitch, dual-row connectors. The internal mechanical structures holding these circuit boards were the result of careful and interactive electrical and mechanical co-design. The end result is that the boards are installed to blind-mate their connectors to their respective downstream board(s), after which they are fastened in place mechanically.

Though it is difficult to render a 3D interconnect topology in two dimensions, Figure 20 illustrates the interconnection of the circuit boards in the palm, as well as the motor control board of a single finger.

3.8. Motor control

Commercial turn-key motor controllers are often designed for a wide range of motors and motor types. As such, they tend to be larger and more expensive than a motor controller designed for a single specific application. For this reason, we chose to develop our own integrated motor control system. To control the fabrication cost of the design, commercial off-the-shelf (COTS) low-cost outrunner brushless motor stators and rotors were disassembled and then re-assembled around the custom motor controllers. Each motor control board contains the necessary hall-effect sensors, amplifiers, and computation and communications resources to drive three motors.

The motor control boards are mounted in between the rotors and the finger heatsink, as shown in Figure 22. This placement allows the hall-effect sensors for brushless commutation to be placed on the same circuit board as the amplifiers, reducing the cost, volume, and reliability issues associated with bringing these sensor inputs to the amplifiers via off-board wiring.

The remainder of the motor-controller wiring shown in Figure 19 was eliminated by soldering the stator windings directly to the motor control boards. Traces on the motor control boards were brought to open pads directly underneath the termination points of the magnet wire phases on the stators, allowing direct and secure connections inside the volume of the enveloping rotors and avoiding any chance of wires being ingested into the assembly.

On the motor control boards themselves, integrated motor drivers were utilized. The ST L6229Q drivers were selected primarily for their compact size and high level of integration, as they contain commutation logic, current control loops, and the three half-bridges required for brushless commutation within a 5 mm $\times$ 5 mm package. Even with the large interior cutouts required in the motor board for the stators to pass through, the resulting layout easily fits within the volume of the finger module, as shown in Figure 22.

However, one penalty for such a high level of power-stage integration is the relatively high bridge resistance of approximately 2 Ω. To manage the associated heat dissipation, the amplifiers are mounted on the bottom side of the motor control board, which is subsequently clamped to an aluminum thermal plate, as shown in Figure 21.

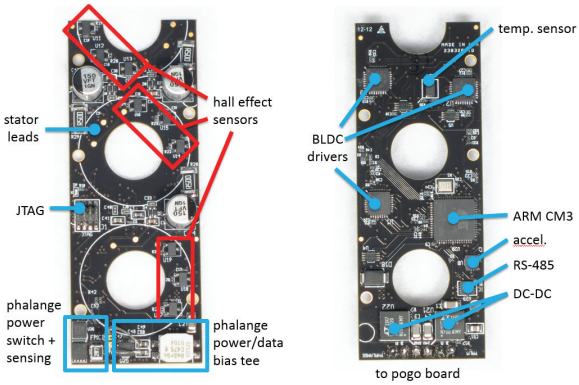

Fig. 22. Finger motor controller board (FMCB).

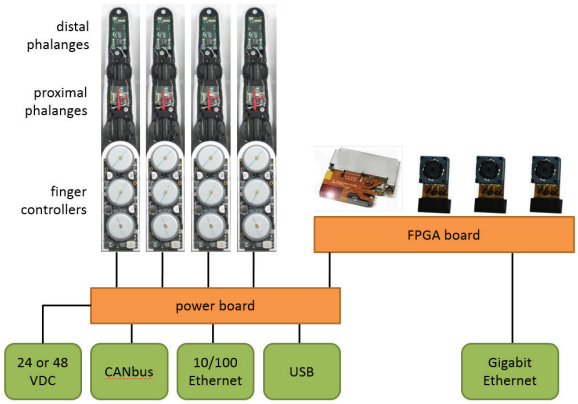

Fig. 23. Data bus topology of the robotic hand.

3.9. *Computational systems*

The hand contains a distributed network of processors: three ARM Cortex-M3 microcontrollers in each finger module, two in the hand frame, and one Xilinx Spartan-6 FPGA in the palm. These computational resources are connected in a hybrid star-like topology, as illustrated in Figure 23. The microcontrollers in the base of each finger module have dedicated point-to-point RS-485 connections to the hand frame. In addition, each finger has an internal phalange power/data bus connecting the three microcontrollers of the finger, as described in Section 3.7. The hand frame microcontroller implements a variety of outside-facing peripherals, such as a CANbus isolator, 10/100 Ethernet, and a USB isolator, as well as an internal SPI link to the FPGA, which in turn offers a high-bandwidth gigabit Ethernet interface. This architecture was designed to allow a variety of connectivity options to upstream hardware in the often-noisy electrical environment of a robotic arm.

Because of the relatively large number of microcontrollers, attention was given to designing methods to allow batch programming of the entire hand. It would be unfortunate, for example, if firmware updates required an operator to manually connect programming adapters to all 12 ARM processors in the fingers. To prevent this situation, a custom bootloader was written that is compatible with the packet structures used during normal operation of the hand. The bootloader waits for several seconds after power-up before booting the application image. Electrical power to the nodes on each data bus can be shut down electronically by its upstream processor. This allows a reprogramming cycle to be triggered at any point by forcing a hard reset of the downstream processors, catching their respective bootloaders, and transmitting a new application flash image. Through this method, a single script on a host computer can reprogram all 12 finger microcontrollers.

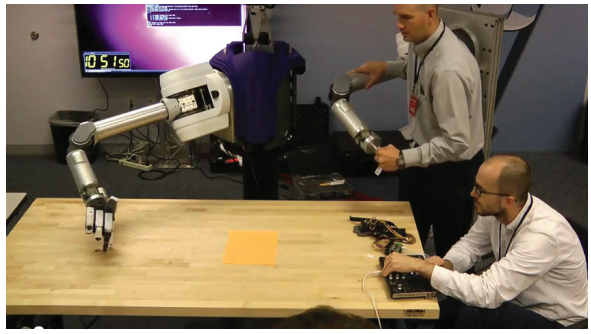

Fig. 24. Experimental setup. One operator controls the slave WAM with a master WAM. Another operator controls the Sandia Hand with either a data glove or a slider panel.

4. Experimental evaluation

The ability of the hand design to grasp and manipulate objects was evaluated by experiments in which the hand attempted to grasp and manipulate a variety of common objects. These experiments also provided feedback about various design decisions, including kinematic topology, actuation strength, thermal management, and skin design. As we will show, the Sandia Hand was capable of grasping or manipulating all of the test objects provided by the experimental procedure.

4.1. *Experimental setup*

The Sandia Hand was mounted to a whole-arm manipulator (WAM, Barrett Technology, Cambridge, MA). Because the objective of the experiment was to evaluate only the hardware, control of the hand was provided solely through teleoperation. Two teleoperators controlled the apparatus: one operator controlled the WAM arm while another teleoperator controlled the hand. Arm teleoperation was achieved by manually moving a master arm which was an identical copy of the WAM, placed near the slave arm with the same orientation. A wand with a spherical tip was attached to the wrist of the master WAM. The arm teleoperator controlled the slave wrist with the wand and controlled the arm posture by moving the master elbow with the other hand, as shown in Figure 24.

We found that some tasks were more readily accomplished by controlling the hand with a data glove, while others were accomplished more easily using a human interface of physical slider controls. We found the data glove to be especially useful for performing high-DOF gross manipulation tasks that required finger gating, but less effective at performing very precise grasping. In experimenting with commercially available data gloves, we found that such gloves tended to exhibit 'mechanical cross-talk' between the flex sensors in the proximal knuckles, as they tended to migrate in between the knuckles of the teleoperator during extended experimental sessions. We therefore designed and built a custom exo-skeletal data glove. While more complex than commercially available data gloves, we found this design to be less prone to these cross-talk issues.

In situations in which highly precise grasp control was required, such as picking up a key from a table, a slider control was mapped to close all fingers simultaneously in a particular grasp type, following the eigengrasp methodology. A particular grasp type was selected and then realized by moving the slider. Cylindrical, spherical, and prismatic grasps (among others) were supported by this interface. Commodity audio-mixing panels were found to have a large number of sliders and are readily available, inexpensive, and easily interfaced via USB.

4.2. Experimental procedure

The slave WAM with the Sandia Hand attached was placed near a flat table top. An orange 8.5 × 11 inch piece of paper was placed in the left half of the table space. One of 19 different objects was placed in an arbitrary position and location on the right half of the table.

The objective was to either grasp the presented object and move it to the orange piece of paper, or perform a combination of manipulation and grasping motions with the object. Table 1 shows the list of objects and the required action to be taken by the hand for each object. Because this test was conducted via teleoperation, neither the vision system nor the tactile sensors were used.

Each object was presented five times and placed in a different location and orientation each time. The grasping and manipulation events were timed. Time began once the hand left a starting position and approached the object on the table. Task completion time was recorded when the object had either undergone the specified manipulation or had been grasped and placed on the orange piece of paper. If grasping could not be completed within 300 s of the start time, the maximum score of 300 s was recorded. For the 15 lb mass, the objective was only to determine if the hand was capable of lifting the mass. As a result, the task was not timed and was performed only once. The measure was either success or failure for that task.

4.3. Results

The task completion times are shown in Figure 25. On two occasions, the battery rolled off the table and out of the workspace of the arm during our first attempts at insertion of the battery into the flashlight. In these instances, the maximum time was recorded. Unlike the grasping tasks, many of the manipulation tasks required the object to be re-oriented and/or re-positioned within the workspace of the arm prior to realizing the manipulation. This part of the task was often the most time-consuming. Example tasks where this was the case include the key, flashlight, hammer, drill, screwdriver, and trowel. The task in which we depressed the push-to-talk button on the two-way radio took the greatest amount of time to accomplish. In three of the five occasions, we were unable to complete the task within the time limit. The Sandia Hand successfully lifted the 15 lb barbell. An array of photos of the Sandia Hand performing these tasks is shown in Figure 26.

At the end of the extensive testing, we touched the motor module covers to determine how hot they were. While they were warm, they were not hot enough that we could not leave our hand on the aluminum covers. The temperature at which it is too painful to maintain contact for an extended period of time is approximately 50°C. We estimate that the aluminum covers were near 40°C, indicating that the design for thermal management was effective at the room temperature of the testing environment.

Table 1. Experimental objects and actions.

Object	Action
ball	grasp
pipe	grasp
canteen	grasp
rock	grasp
water bottle	grasp
telephone	grasp
floodlight bulb	grasp
pelican case	grasp
15 lb barbell	grasp
key	large manipulation (re-orient in stable grasp)
flashlight	internal articulation (depress trigger)
hammer	large manipulation (finger gate into stable grasp)
drill	internal articulation (depress trigger)
radio	internal articulation (depress push-to-talk switch)
screwdriver	large manipulation (finger gate into stable grasp)
trowel	large manipulation (finger gate into stable grasp)
wire cutters	internal articulation (actuate to cut wire)
AA battery	grasp
1/4 turn valve	internal articulation (turn valve)

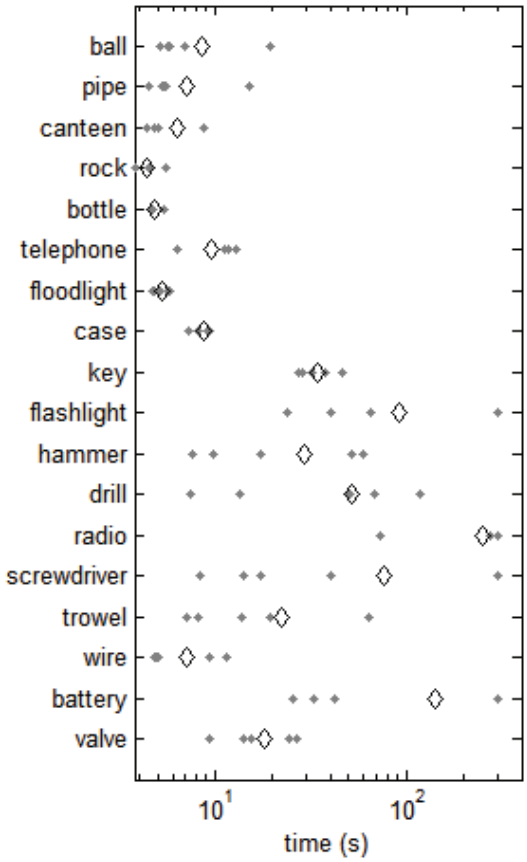

Fig. 25. Task completion test results. Individual trial results are shown with gray dots. Averages across those trials are shown with a diamond.

Fig. 26. Photos of the Sandia Hand performing the experimental tasks.

5. Conclusions and future work

A robotic hand was designed, prototyped through numerous iterations, and shown to be capable of performing a variety of grasping and manipulation tasks. Design tradeoffs were made to balance the performance objectives of manipulation capacity, grasp quality, sensing, modularity, robustness, volume, mass, and cost. Manipulation capacity was afforded by four 3-DOF fingers. Sensing was implemented using low-cost mass-market sensors for position, tactile, and inertial sensing. A modular design was implemented at the finger level with finger modules, as well as at the finger subsystem level with easily separable subsystems. Robustness was provided by sound structural design as well as mechanical fuses between the hand frame and the finger modules. Volume was reduced by selecting efficient and compact transmission topologies, and eliminating electrical wires. Mass was reduced by selecting drive components (often the heaviest components) which have high power and torque density. Finally, cost was reduced by selecting only mechanical and electrical components that exist in mass-market products, and thus benefit from economies of scale without requiring the robot hand itself to be manufactured in similarly large quantities.

Future work will include explorations of alternative kinematic topologies of the finger modules. While the presented topology was effective, particularly in the teleoperative case, there may be better solutions. There are a number of grasping analysis tools that could be used to guide future arrangements of finger modules. To our knowledge, manipulation analysis tools are far less advanced. Research in the area of robot-specific simulation of manipulations would contribute significantly to the development of robot hands by offering quantitative guidance to kinematic designers.

Future work will also evaluate the utility of our tactile sensors and hand-based vision sensors. The quantitative evaluation presented in this paper was performed by teleoperation. As a result, the utility of the tactile and vision sensors was not evaluated. Autonomous operation of the hand, however, will rely heavily on these sensors and would provide important insight into their utility and areas for future improvement.

We are encouraged by the evaluation results, which suggest that the design concepts presented in this paper may

have commercial value. Enabling the distribution of dexterous hands to more researchers at a lower cost could subsequently enable the development of software capable of performing dexterous grasping and in-hand manipulations. The result will be an accelerated pace of the development of high-DOF grasping and manipulation software, allowing robotic hand technology to contribute to a variety of viable commercial applications.

Acknowledgments

This work was completed while Morgan Quigley was a PhD candidate at Stanford University and Curt Salisbury was with Sandia National Laboratories, Albuquerque, NM, USA.

Funding

This work was funded in part by the DARPA Autonomous Robotic Manipulation (ARM) Program.

References

Barbagli F, Prattichizzo D and Salisbury J (2005) Modelling and controlling the compliance of a robotic hand with soft finger pads. In: *Proceedings of the first workshop on multi-point interaction with real and virtual objects*, pp. 125–137.

Bicchi A (2000) Hands for dexterous manipulation and robust grasping: A difficult road toward simplicity. *IEEE Transactions on Robotics and Automation* 16(6): 652–662.

Bicchi A and Sorrentino R (1995) Dexterous manipulation through rolling. In: *Proceedings of the 1995 International conference on robotics and automation*, pp. 452–457.

Birglen L, Laliberte T and Gosselin C (2000) *Underactuated Robotic Hands*. New York, NY: Springer.

Butterfass J, Hirzinger G, Knoch S, et al. (1998) DLR's multisensory articulated hand part I: Hard- and software architecture. In: *Proceedings of the IEEE International conference on robotics and automation*, pp. 2081–2086.

Higashimori M, Jeong H, Ishii I, et al. (2005) A new four-fingered robot hand with dual turning mechanism. In: *Proceedings of the IEEE International conference on robotics and automation*, pp. 2679–2684.

Jacobsen S, Iversen E, Knutti D, et al. (1986) Design of the Utah/M.I.T. dextrous hand. In: *Proceedings of the IEEE International conference on robotics and automation*, pp. 1520–1532.

Konolige K (2010) Projected texture stereo. In: *Proceedings of the IEEE International conference on robotics and automation*, pp. 145–155.

Lim J (2009) Optimized projection pattern supplementing stereo systems. In: *Proceedings of the IEEE International conference on robotics and automation*, pp. 2823–2829.

Lincoln LS, Quigley M, Rohrer B, et al. (2012) An optical 3D force sensor for biomedical devices. In: *Proceedings of the International conference on biomedical robotics and biomechatronics*, pp. 1500–1505.

Markenscoff X, Ni L and Papadimitriou CH (1990) The geometry of grasping. *The International Journal of Robotics Research* 9(1): 61–74.

Mason M and Salisbury J (1985) *Robot Hands and the Mechanics of Manipulation*. Cambridge, MA: MIT Press.

Salisbury J and Craig J (1982) Articulated hands: Force control and kinematic issues. *The International Journal of Robotics Research* 1(1): 4–17.