Wolverine: A Wearable Haptic Interface for Grasping in Virtual Reality

Inrak Choi, Elliot W. Hawkes, David L. Christensen, Christopher J. Ploch, and Sean Follmer[1]

***Abstract*— The Wolverine is a mobile, wearable haptic device designed for simulating the grasping of rigid objects in a virtual reality interface. In contrast to prior work on wearable force feedback gloves, we focus on creating a low cost and lightweight device that renders a force directly between the thumb and three fingers to simulate objects held in pad opposition (precision) type grasps. Leveraging low-power brake-based locking sliders, the system can withstand over 100N of force between each finger and the thumb, and only consumes 0.24 mWh (0.87 joules) for each braking interaction. Integrated sensors are used both for feedback control and user input: time-of-flight sensors provide the position of each finger and an IMU provides overall orientation tracking. This paper describes the mechanical design, control strategy, and performance analysis of the Wolverine system and provides a comparison with several existing wearable haptic devices.**

I. Introduction

Though Virtual Reality (VR) has been explored in research contexts since the late 1950s, recent advances in display technology have made consumer VR a reality. While new devices such as the Oculus Rift or HTC Vive provide high resolution visuals, the user input devices have been limited to traditional game controllers and existing styles of gestural input. It is desirable to allow users to touch what they can see and physically manipulate virtual objects. However, current consumer input devices do not provide the kinesethic feedback that we experience when interacting with objects in the real world.

Ideally, haptic feedback interfaces for consumer VR should be low cost, lightweight, ungrounded, while still providing force feedback that realistically simulates touching and manipulating objects; that is, the interfaces should resist forces larger than finger strength at a high refresh rate with high accuracy. There have been a variety of approaches in the research literature that explore force feedback gloves: externally grounded systems [1], [2], systems grounded to the wrist [3], systems providing forces between the palm and fingers [4], and systems providing forces between the thumb and fingers [5]. However, none of these devices meet all of the above design objectives for consumer based devices. We seek to meet these design objectives for the specific case of grasping virtual rigid objects in precision-based grips.

Accordingly, we present the design of a device, termed the Wolverine system, that attaches to the tips of three fingers and the thumb (Fig. 1). The device provides the sensation of grasping a rigid object by resisting relative motion between the fingers and thumb. It utilizes a brake-based system to provide high resistance to forces in a light-weight, low power, and low cost package. In this paper we describe the Wolverine system's mechanical design, integrated sensing and control, and provide an analysis of its performance.

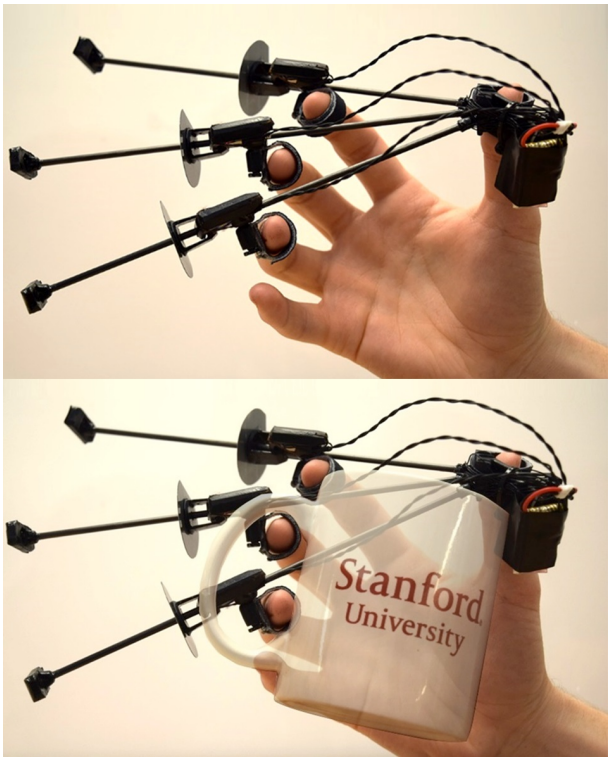

Fig. 1. Wolverine, a new wearable haptic user interface for grasping in virtual reality, holding a cylinder-shaped virtual object.

II. Related Work

Researchers have developed externally-grounded haptic interfaces with external actuation in order to make the manipulator light. The most popular haptic interface in this category is the PHANToM [1] which allows a user to feel stiffness and textures of virtual objects through active force feedback on a finger mounted end effector. HIRO [6] and SPIDAR [2] can also be included in this category. While previous research has investigated brake-based haptic interfaces, they primarily use passive force feedback with brakes for guidance in path following applications tasks[7], [8], [9]. These grounded haptic interfaces create various types of force feedback since the device is fixed on the ground, but the work envelope is limited to a small 3d space.

Researchers have also developed glove-style haptic interfaces to give users more degrees of freedom in motion. The first glove-style haptic interface, CyberGrasp [3], was launched commercially in early 1990s. Since then, there have

[1]The Department of Mechanical Engineering at Stanford University, 450 Serra Mall, Stanford, CA 94305, USA irchoi@stanford.edu

been other exoskeleton force-feedback gloves developed using different mechanisms, such as passive spring and clutch force feedback devices [10], wire-driven devices [11], [12], magnetorheological fluid devices [13], [14], [15], and micro hydraulic systems [16]. These systems are grounded to back of the hand or wrist of the user. One limitation of such systems could be unexpected kinesthetic or tactile feedback at the contact area between the device and user's hand.

Other researchers have explored providing force directly between the fingers and the palm to simulate palm opposition type grasping, such as the Rutgers Master II (RMII) [4]. More recently, researchers have investigated the use of particle jamming to provide resistance between the fingers and palm [17], [18]. Most related to our work are devices which provide forces directly between the fingers and thumb to simulate pad opposition or precision type grips. Zhang et al. explored lightweight electroactive polymer actuators between the thumb and forefingers (DESR), however it has a limited range of motion [5]. Our focus is on supporting a wide range of motion in a lightweight, low-cost package; however, in order to achieve this goal, we sacrifice active force feedback and the ability to render variable stiffness.

III. Design

A. Overall Structure

The Wolverine is composed of a base, which mounts on the thumb, and three connected rods, each of which has a sliding mount for the tips of the index, middle, and ring fingers (Fig. 2). Each sliding mount has a brake that can lock onto the respective rod. Therefore, the three finger tips are physically connected to the thumb tip through an exoskeleton structure that can generate precision grasping motions [19], [20]. The rods are connected to the base with ball joints (3 DOF each), and the sliding mounts are connected to the rods with cylindrical joints (2 DOF each). The supporting structures physically in contact with the three finger tips are connected to the sliding mounts with revolute joints (1 DOF each). Due to its many degrees of freedom and low friction and inertia, this structure allows the hand to move freely. However, when it is desirable to create the feeling of grasping an object, brakes are actuated that lock the sliding mounts at desired locations on the rods. As a result, the kinesthetic force feedback that would be felt when gripping a rigid object is recreated.

Carbon fiber tubes are used for the three rods and rest of the mechanical components are printed on a SLA 3D printer. Each brake module including a dc motor weighs only 5 grams, and the total weight of the device sums up to be approximately 55 grams, including a 350 mAh battery. Power is transmitted attached through an electric wire from the thumb mounted control board. The device is wireless and communicates with a master controller through a Bluetooth module. Thus, external cables do not impede arm motions and reduce realism in virtual reality applications.

Fig. 2. Wolverine system overview: *Top Right* shows the sensing degrees of freedom, *Center* shows the motion degrees of freedom, and *Bottom* shows closeup views of the individual components.

B. Actuation

We use a brake mechanism to render virtual objects in order to make the Wolverine compact and energy efficient; the previously mentioned haptic gloves have chosen active actuators for force feedback to generate variable stiffness. A brake system, in general, guarantees stable motions because it can only dissipate energy, and is often more compact than active actuators of the same strength. However, brake systems can only resist motion, which could lead to an unnatural grasping sensation if there is any resistance when the user opens his or her hand.

Therefore, we propose a mechanism for directional braking in haptic applications. As shown in Fig.3 (A), in the default state with the actuation off, the levers is in the "out" position. The hole in the lever through which the rod passes is coaxial with the rod, allowing the sliding mount to move freely along the rod. In order to lock the sliding mount with respect to the rod, a 6mm diameter 20:1 geared dc motor is turned ON and pulls a wire that rotates the lever counterclockwise taking roughly 20 ms. Now the hole through which the rod passes is no longer coaxial with the rod, and jamming occurs. As the user applies a force pulling the sliding mount toward the thumb, Tendon 1 becomes taut, further rotating the lever counterclockwise. The controller

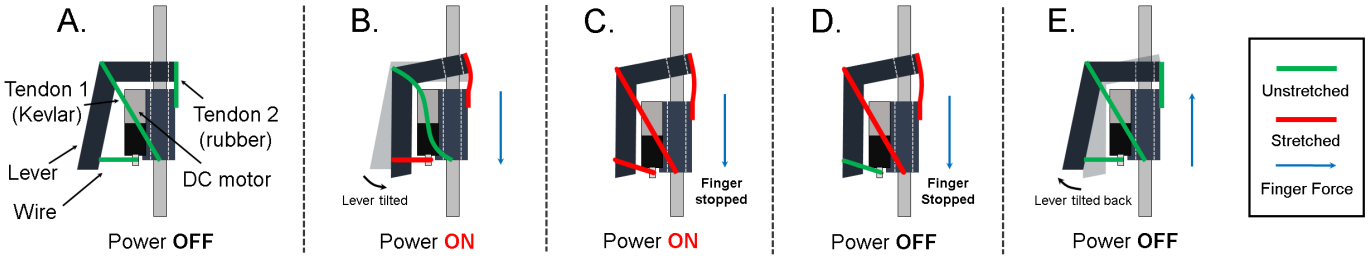

Fig. 3. Brake mechanism. A) When the power is OFF, the lever is in the "out" position, and the hole in the lever through which the rod passes is aligned with the rod. B) When the motor is powered ON, a wire is pulled, rotating the lever to the "in" position; the rod now jams in the hole in the lever. C) As a load is applied by the user to the sliding mount, Tendon 1 becomes taut, further jamming the rod in the hole in the lever. D) Even when the power is turned OFF, the brake remains engaged due to tension in Tendon 1. E) When the user releases the squeezing force, the elastic Tendon 2 pulls the lever back into the "out" position, and the sliding mount is free to move away from the thumb.

then turns off the power, but the brake is still engaged. This stage can last as long as necessary to complete the desired task in virtual reality. Once the user finishes the task and opens his or her hand, the elastic Tendon 2 rotates the lever back clockwise, unlocking the braking mechanism. The user can freely move his or her finger away from the thumb.

It is important to note that the dc motors are only used for initiating braking (Fig.3 (B-C)), but not required to maintain braking once the user is applying a force that pulls the sliding mount toward the thumb (Fig.3 (D)).

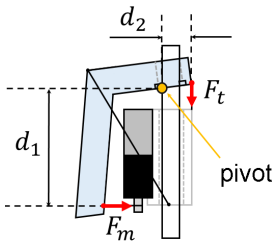

Fig. 4. Free body diagram of initial brake engagement

1) Active Brake Engagement: The active engagement of the brake is the heart of the device, governing both actuation speed and the accuracy of the output. It is therefore helpful to introduce the parameters that determine its performance.

Fig. 4 shows a simple free body diagram of the brake lever, pivoting about its contact point. The motor is the active element that engages the system, while Tendon 2 (labeled in Fig. 3) provides the force that keeps the brake from jamming unintentionally, and in the current prototype also passively back-drives the motor for release. For the sake of analysis, consider the static case of moment balances about the pivot:

$$\sum M = d_1 F_m - d_2 F_t = 0 \tag{1}$$

where d_1 is the height of the lever, F_m is the force of the motor, d_2 is the distance from the pivot to the spring and F_t is the force of the spring.

The force from the tendon spring (F_t) and the force from the motor (F_m) are modeled by:

$$F_t = k_t * (l_f - l_i) \quad \& \quad F_m = T_m / r_m \tag{2}$$

where k_t is the spring constant, and l_f and l_i are the final and initial tendon spring lengths respectively, T_m is the motor drive torque and r_m is the motor pulley radius.

Using Eq.1 and Eq.2, we find the minimum motor torque necessary to engage the brake:

$$T_{m\_min} = d_2 r_m k (l_f - l_i)/d_1 \tag{3}$$

The minimum tendon spring tension required to back-drive the motor is governed by:

$$F_{t\_min} = r_m T_{backdrive} \tag{4}$$

Where $T_{backdrive}$ is the torque required to backdrive the motor.

We note that since, the forward motor torque depends on the spring force, and the spring force depends on motor back-drivability, the ratio of forward motor torque to backdrivable motor torque is important. This is part of the reason we choose a motor with a relatively small gear ratio (20:1) in the current prototype. Actively reversing the motor for a short duration would eliminate this issue.

C. Sensing

Sensors are integrated into the system in order to measure the linear position of each finger tip along the rod as well as the overall orientation of the device and hand. Specifically, a Time-of-Flight (ToF) sensor (STMicroelectronics VL6180X) is mounted at the tip of each rod and measures the time that emitted IR light takes to travel to a reflective pad on the sliding mount and return to the sensor. Because we measure position between the finger and the end of the rod rather than between the finger and the thumb, the sensor does not interfere when the finger and thumb are brought close to one another. Optical sensing generally is attractive to reduce weight and inertia of moving parts and adds no friction to the system. Within the realm of optical sensing, ToF sensing has the benefit that the signal is relatively insensitive to ambient light conditions and the quality of the reflective pad or how it varies over time (for example with smudges or dust). Linear variable differential transformer (LVDT) sensors, or linear capacitive sensors like those used in digital calipers would also be convenient for this form factor, but would require more systems integration work.

A 9 axis inertial measurement unit (InvenSense MPU9250 with sensor fusion) is coupled to a thumb to measure orientation. With a single orientation sensor, we assume the thumb represents the orientation of the hand. However, an IMU could be added to each finger in future versions if user studies show that the pose of individual fingers is important information.

D. Control

A number of features of the design result in simple control. Measurement and modeling of the hand's pose is simple because the device directly measures the distances between the fingers and the thumb. This is in contrast to devices that measure joint angles. In such devices, computational effort is needed to model fingertip motion by forward kinematics, and there is a possibility of accumulating error from joint to joint.

While other haptic gloves focus on generating realistic stimuli of soft objects [21], [22], our device is a position-control rather than force-control device, and therefore only renders rigid bodies. While limiting the system to the display of rigid bodies does reduce its capabilities, many objects in our daily lives can be approximated as rigid.

With the choice of only reproducing rigid objects comes greatly simplified computation. Computing rigid contact location using a simple boundary is much less computationally expensive than soft interactions or simulated rigid contact through real-time force rendering using finite element analysis. Further, the use of friction as the force generation method results in a passively stable system without the need for active control once the brake is engaged.

IV. Performance Analysis

A. Actuation Speed

To decide the voltage for actuation, we powered the dc motor with different voltages. As shown in Fig.5, higher voltages made faster responses. Also, the variance of the actuation lag was reduced in higher voltages. This means a higher voltage can generate faster and more reliable motions for this system. However, repeated operation at 5V resulted in damage to the gearbox. Therefore, for this system 3.7 volts was chosen.

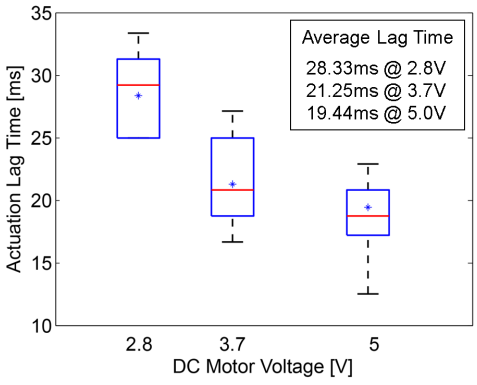

Fig. 5. Time for the dc motor to tilt and engage the locking levers at various voltages.

With 3.7 volts, the average time to rotate the lever into the locked, or "in," position is 21ms. This actuation speed is less than the delay a human notices between visual and haptic stimuli (45ms) [23], [24]. The actuation time could be further reduced by decreasing the angle through which the lever must rotate. However, this could decrease robustness, because the system is closer to jamming in the unlocked position; a small perturbation could result in undesired locking.

B. Force Analysis

To characterize its stiffness and maximum force, we mounted a set of actuation parts to Instron MicroTester 5848 and measured the compression force with the brake engaged. The force-displacement curve in Fig.6 shows the stiffness during braking is 162N/mm. The maximum force before slipping is 106N. The brake force is a frictional force between the carbon fiber rod and an aluminum flat washer affixed to the lever.

The current force level of over 100N is suitable for our application because it is larger than the forces generated from precision grasps. Previous work in grasping describe the average strength of chuck pinch as 7.9kg (77.4N) for men, and 5.2kg (51.0 N) for women [25]. While other devices are able to render more complex force profiles, their maximum forces are substantially lower ranging from 5N - 29N [16], [14], [5], [26], [4], [11], making them more suitable to delicate tasks.

The device's stiffness and maximum force could be increased further by changing the materials of the device. Currently, we are using carbon fiber tubes for the rods and a liquid resin material for other parts. Using metals instead would increase the stiffness and maximum force.

C. Sensor Noise Analysis

1) Noise dependence on static displacement: While a ToF sensor's sensitivity should be unaffected by the size of the reflecting surface (within the field of view) and its surface properties, the magnitude of the signal is affected. Therefore, the noise performance of the sensor depends on the design of the surface reflector. Fig. 7 shows the one standard deviation

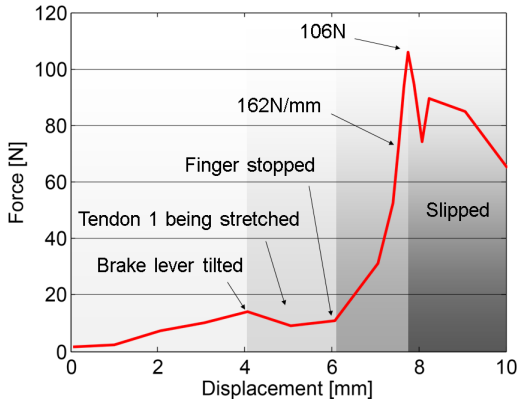

Fig. 6. Force - displacement curve of the brake mechanism. Measured by Instron 5848.

noise magnitude as a function of displacement from the sensor to the back of the finger for 3 different reflector designs. We note that for small distances, less than roughly 50mm, all reflectors are fully within the field of view and so have comparable noise performance–likely determined by characteristics of the sensor and the material of the reflector. However, after a displacement of about 50mm, the noise begins to diverge for the smaller disks. The 37mm disk size was found to be the largest size possible that did not interfere with the grasp mechanics.

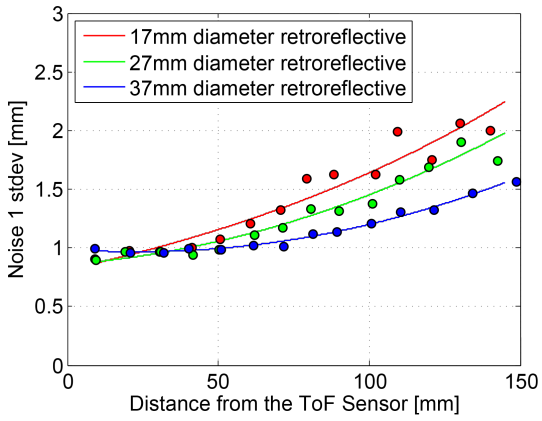

Fig. 7. Relationship between the ToF sensor noise and the distance from the sensor to the reflector, mounted on the back of a finger. Data taken at 50Hz sampling rate and plotted with corresponding second order fits.

It is not ideal that the noise increases with larger distances between the sensor and reflector, because a large distance occurs when a user grasps a small virtual object. Users are more sensitive to error grasping smaller objects. However, the magnitude of the noise (1.5mm) is roughly the same as the just noticeable difference (JND) for humans performing small grasping tasks [27], [26], [28]. While inverting the system (sensing the distance from the fingers to the thumb directly) would decrease noise for small virtual objects, this choice would limit the reflector size due to geometric constraints and would perform much worse for all but the smallest grasps. If less noise is desired, an LVDT sensor could be implemented.

2) Resolution versus Sampling Rate: The ToF sensor performs many fast measurements of distance and averages them down to the output frequency (100Hz) to reduce noise. For operation at slow grasping speeds, we can similarly take a running average of the resulting data to get sub-millimeter position estimates. This utilizes a measurement rate after averaging that is lower than the 100Hz sensor sampling rate. Since the system only outputs integer values corresponding to millimeter length, a slow grasp will find many repeated measurement values between whole millimeter length transitions. By averaging these and assuming constant velocity over the course of the averaging time, one can achieve sub-millimeter estimations of position. This, however, comes at the expense of temporal resolution. An ideal system would dynamically scale the overall sensor bandwidth according to grasping speed to give increased spacial measurement accuracy for low speed grips, and increased temporal accuracy at high speeds.

In practice, we cannot assume a constant velocity of movement by the user. This means that continually decreasing the measurement rate after averaging does not necessarily result in decreased uncertainty. To illustrate the relationship between measurement rate after averaging and position uncertainty with non-constant velocity, we assume that the non-constant velocity results in a temporal uncertainty of roughly 10% of the total time period (the actual uncertainty will be investigated in future human studies). We plot the simulated results in Fig. 8. The position uncertainty at a low measurement rate after averaging is very large because the measurement is not keeping up with the real motion. At a high measurement rate after averaging, we see an increase in noise in the measurement due to the increased sensor bandwidth. For slow grasps we see a minimum in uncertainty where these competing effects balance. Given the slopes on both sides of this minimum it is safer to favor the side of high measurement rate after averaging. Further, for quick grasps in this range, a higher measurement rate after averaging results in lower uncertainty, monotonically.

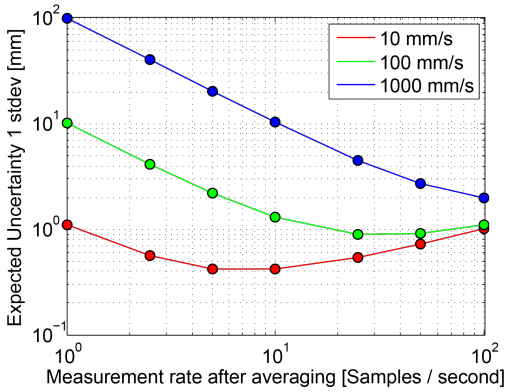

Fig. 8. Simulation of expected position uncertainty vs measurement rate after averaging for a variety of grasp speeds (10-1000mm/s). This assumes a temporal uncertainty is of 10% of the effective measurement period. At low grasp speeds, a low measurement rate after averaging can result in lower uncertainty than a higher measurement rate after averaging.

3) Spatial Uncertainty: There are 4 main sources of spatial uncertainty for the device attempting to grip at a known distance:

a) *Sensor Noise* is shown to vary inversely with grip size (Sec. IV-C.1) and can be made through further averaging (Sec. IV-C.2) to decrease with slower grip speeds. Noise is unpredictable in nature and cannot be removed with modeling.

b) *Sensor Lag* is important during fast grasping motions. With the current sampling rate of 100 Hz sampling rate the system only gets an updated position every 10ms, meaning the data is, on average, 5ms old. The uncertainty due to this lag will increase linearly with grasping speed, and can be removed with a forward model (Fig. 10) assuming constant velocity.

c) *Actuator Lag* has been found to be about 20ms on average (see Fig. 5). Like sensor lag, the resulting error will be linear with grasping speed, and can be removed with a forward model as well (Fig. 10).

d) *Actuator Lag Variance*, like actuator lag, will result in an error that is linear with grip velocity. However, like sensor noise, this uncertainty is not predictable because is represents *variability* in the actuation time itself. As can be seen in Fig. 5, the faster average actuation times result in smaller variances reducing the impact of this term.

Sensor noise(a) and sensor lag model errors(b) are considered to be independent, as are actuator lag variance(c) and actuator lag model errors(d). This means that at least to the first order these can be all treated as independent uncertainties and so should add according to:

$$\sigma_{total} = \sqrt{{\sigma_a}^2 + {\sigma_b}^2 + {\sigma_c}^2 + {\sigma_d}^2} \tag{5}$$

We note that the specific details of error and whether it occurs before or after forward modeling could cause coupling between the uncertainty terms in application due to a possible shared dependency on velocity.

When properly filtered, sensor noise should be proportional to velocity. The remaining error in sensor lag and actuator lag after feed forward canceling (due to the imperfect constant velocity approximation) should also be proportional to velocity (but is likely small), as should actuator lag variance. This results in our total uncertainty being proportional to grip velocity, and having a magnitude of roughly 10mm for the fastest grip rates.

D. Power Consumption

We use a 3.7V 350mAh battery to run the device. The electrical current consumption was measured using a power supply to simulate the battery at 3.7 volts. The baseline current was found to be about 50mA for the microprocessor, finger sensors, and thumb orientation sensor. When the finger brake motors are actuated, the total current increases to 780mA for the 320ms actuation cycle. This design enables us to use the device for roughly five hours with 1500 full fingered grasping events.

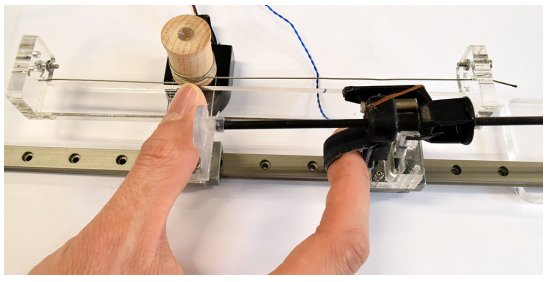

Fig. 9. Test rig for measuring absolute distance and velocity of the device using an encoder in addition to the onboard ToF sensor.

V. Results and Discussion

A. Uncertainty Verification

To analyze errors in actual locking positions, we mounted the device to a linear guide with an incremental rotary encoder as shown in Fig.9. A capstan drive is used to connect the encoder shaft and device with minimal friction. The encoder's 1024 cycles per revolution in quadrature result in a linear resolution of 14.2 μm. Velocity is calculated at 1000Hz.

We ran two tests, first analyzing the uncertainty due only to actuation lag and how well the implemented actuation lag compensation works, and second analyzing the uncertainty due to both the actuation and sensing/communication lag and how well the full (actuation plus sensing/communication) compensation works. Such compensation helps improve performance, especially at high grasping speeds, where expected error is over 20mm.

1) Actuation Lag Compensation: To analyze the actuation lag error alone, we set the desired position to 50mm while sensing with the encoder for control, not with the on-board ToF sensor. This allowed us to temporarily remove errors from the on-board position sensor to accurately characterize errors due to actuator lag alone.

The results are shown in Fig.10 as the light blue points, and together represent an actuator lag of 23.4ms which is similar to the values previously shown in Fig.5. This average lag time generates a distance offset that increases linearly with grasping speed. We also note a small constant distance offset of 2.9mm, regardless of grasping speed, due to a small amount of backlash in the system. These experimental values allow us to build a forward model to compensate both the distance offset due to actuation lag and the constant offset. The equation governing this is given by:

$$d_{act\_model} = d_{des} + v * t_{a\_lag} + d_{const} \tag{6}$$

where d_{des} = 50mm, t_{act_lag} = 23.4ms, d_{const} = 2.9mm

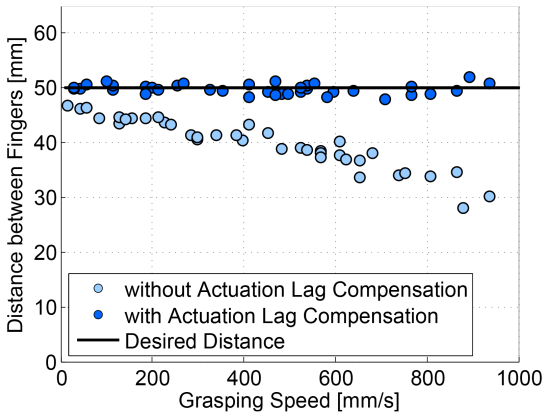

Fig. 10. Test results showing the distance between fingers at various grasping speeds, using an encoder to measure distance; therefore only uncertainty due to actuator lag is present (no sensing/communication lag). Light blue points are data with no forward compensation. Dark blue points are the results after forward model compensation.

TABLE I

COMPARISONS WITH OTHER DEVICES.

	Wolverine	CyberGrasp [3]	Rutgers Master II [4]	DESR [5]
Grasp Type	pad opposition only	pad, palm, side opposition	palm opposition only	pad opposition only
Actuator Type	one-way brake	dc motor with wire driven	pneumatic cylinder	electroactive polymer
Force Feedback	constant stiffness	variable stiffness	variable stiffness	variable stiffness
Maximum Force	106N	12N	16N	7.2N
Motion Range	20-160mm	full hand closing	27mm stroke	5mm stroke
Power Source	built-in battery	external cables	external tubes	external cables
Weight	55g	450g	185g	38g

The dark blue points in Fig.10 are the experimental data with this compensation. We see the system now shows uncertainty that is independent of grasping speed. There are still, however, small errors from the desired distance due to imperfect actuation repeatability. This can potentially be improved by adopting a faster or more consistent actuator.

2) Sensing and Communication Lag: To test the uncertainty due to both the actuation lag and the sensing/communication lag, we test the device using the onboard ToF sensor while measuring the actual distance with the encoder. As shown in Fig.11, the error is substantial, especially for high grasping speeds. We add lag due to sensing and communication to the forward model to help mitigate this error. The dominant lag comes from the sampling rate itself; the sampling period is 10ms, resulting in an effective sample lag of 5ms. Further, the I^2C communication takes an additional 0.2ms. These lags can be added to the model:

$$d_{model} = d_{des} + v * (t_{a\_lag} + t_{s\_lag} + t_{c\_lag}) + d_{const} \quad (7)$$

Both these forward models rely on the assumption that our fingers only move in one direction during grasping motions.

Implementing the full forward model for actuation plus sensing/communication lag, we see improved results in Fig. 11.

We see that both the bias and the systematic error with increasing velocity are almost entirely removed. All that remains is the uncertainty that increases with velocity as previously predicted and described. This error is likely due to sensor noise and errors in the assumption of constant velocity because of the difference in errors seen in Fig. 10 (dark blue) and Fig. 11 (red). Actuator lag uncertainty is also present but appears to be less important, especially at higher speeds.

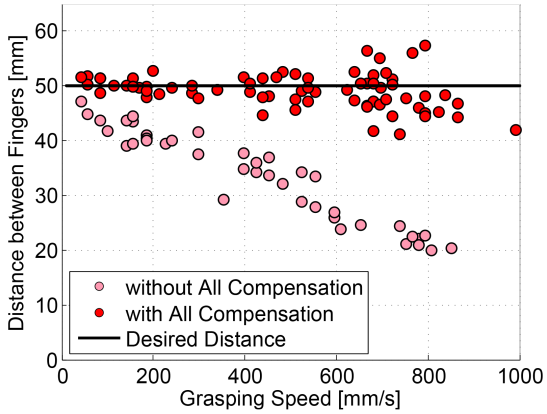

Fig. 11. Test results showing the distance between fingers at various grasping speeds, using the onboard ToF sensor. Uncertainty due to actuator and sensing/communication lag is present. Pink points are data with no forward compensation. Red points are the results after forward model compensation.

B. Possible Virtual Objects

With the wide range of motion (20-160mm) and high stiffness (162N/mm), many objects in our daily lives can be simulated by the Wolverine. It enables human to grasp using precision grips any rigid object larger than a 20mm diameter sphere and not exceeding a 160mm sphere. Based on the YCB Object Set [29], a reference of objects of daily life, about 75% items in the set could be rendered by the Wolverine.

C. Comparison of the Wolverine with Other Devices

Table 1 provides details for comparison with CyberGrasp [3] and other devices. As shown in the table, the Wolverine system can provide a large range of motion and high resistance forces, but all other systems provide variable stiffness. The trend toward mobile VR applications like the Samsung Galaxy VR makes the Wolverine particularly interesting. The lightweight, battery-powered design could be consumer friendly, and the processing of binary output signals is computationally simple, making it feasible even on mobile processors.

VI. CONCLUSIONS AND FUTURE WORK

In conclusion, we have introduced the Wolverine haptic feedback device for virtual grasping of rigid objects. The brake mechanism can provide over 100N of force between each finger and the thumb. With our described forward model and integrated time of flight sensors we can render distances between the thumb and finger with a resolution close to that of human perception. The Wolverine system is lightweight (under 55 grams including all sensors and battery), low power (can run on a 350mAh battery for 5hrs), low cost (under $40 in parts for mass production), and has a large motion range, making it ideal for mobile consumer use. Its major limitation is that it does not render variable stiffness and is therefore suited for only certain applications.

Future work will focus on building up the rest of the infrastructure for performing user studies to validate its use.

For such studies, we would like to integrate the device with a consumer VR display. In order to do so, we need to know the position of all fingers and the thumb. Since finger position is measured locally with respect to the thumb, and the IMU provides the orientation of the thumb, all that is needed is the 3-axis global coordinates of the thumb. For this a simple motion tracking sensor would be adequate. In addition, we plan to characterize the effect of dynamic filtering based on the current grasp speed (to reduce noise), as well as explore faster actuation of our brake mechanism and more accurate and high speed position sensing. All of these changes could potentially improve the accuracy of the haptic rendering.

ACKNOWLEDGMENTS

We thank Mark Cutkosky and the members of the BDML for insightful conversations.

REFERENCES

[1] T. H. Massie and J. K. Salisbury, "The phantom haptic interface: A device for probing virtual objects," in *Proceedings of the ASME winter annual meeting, symposium on haptic interfaces for virtual environment and teleoperator systems*, vol. 55, no. 1. Chicago, IL, 1994, pp. 295–300.

[2] M. Sato, "Development of string-based force display: Spidar," in *8th International Conference on Virtual Systems and Multimedia*. Citeseer, 2002.

[3] "Cybergrasp, cyberglove systems inc." http://www.cyberglovesystems.com/cybergrasp/, accessed: 2016-02-28.

[4] M. Bouzit, G. Burdea, G. Popescu, and R. Boian, "The rutgers master ii-new design force-feedback glove," *Mechatronics, IEEE/ASME Transactions on*, vol. 7, no. 2, pp. 256–263, 2002.

[5] R. Zhang, A. Kunz, P. Lochmatter, and G. Kovacs, "Dielectric elastomer spring roll actuators for a portable force feedback device," in *Haptic Interfaces for Virtual Environment and Teleoperator Systems, 2006 14th Symposium on*. IEEE, 2006, pp. 347–353.

[6] T. Endo, H. Kawasaki, T. Mouri, Y. Ishigure, H. Shimomura, M. Matsumura, and K. Koketsu, "Five-fingered haptic interface robot: Hiro iii," *Haptics, IEEE Transactions on*, vol. 4, no. 1, pp. 14–27, 2011.

[7] D. Gao and W. J. Book, "Steerability in planar dissipative passive robots," *The International Journal of Robotics Research*, 2009.

[8] Y. Matsuoka and B. Townsend, "Design of life-size haptic environments," in *Experimental Robotics VII*. Springer, 2001, pp. 461–470.

[9] Y. Hirata, Y. Tozaki, and K. Kosuge, "Wire-type human support system controlled by servo brakes," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 3356–3361.

[10] T. Koyama, I. Yamano, K. Takemura, and T. Maeno, "Multi-fingered exoskeleton haptic device using passive force feedback for dexterous teleoperation," in *Intelligent Robots and Systems, 2002. IEEE/RSJ International Conference on*, vol. 3. IEEE, 2002, pp. 2905–2910.

[11] K. Koyanagi, Y. Fujii, and J. Furusho, "Development of vr-stef system with force display glove system," in *Proceedings of the 2005 international conference on Augmented tele-existence*. ACM, 2005, pp. 91–97.

[12] L. Jiang, "Portable haptic feedback for training and rehabilitation," p. 114, 2009, copyright - Database copyright ProQuest LLC; ProQuest does not claim copyright in the individual underlying works; Last updated - 2016-05-28. [Online]. Available: http://search.proquest.com/docview/305009113?accountid=14026

[13] S. H. Winter and M. Bouzit, "Use of magnetorheological fluid in a force feedback glove," *Neural Systems and Rehabilitation Engineering, IEEE Transactions on*, vol. 15, no. 1, pp. 2–8, 2007.

[14] J. Blake and H. B. Gurocak, "Haptic glove with mr brakes for virtual reality," *Mechatronics, IEEE/ASME Transactions on*, vol. 14, no. 5, pp. 606–615, 2009.

[15] Y. Nam, M. Park, and R. Yamane, "Smart glove: hand master using magnetorheological fluid actuators," in *International Workshop and Conference on Photonics and Nanotechnology 2007*. International Society for Optics and Photonics, 2007, pp. 679 434–679 434.

[16] Y. Lee and D. Ryu, "Wearable haptic glove using micro hydraulic system for control of construction robot system with vr environment," in *Multisensor Fusion and Integration for Intelligent Systems, 2008. MFI 2008. IEEE International Conference on*. IEEE, 2008, pp. 638–643.

[17] I. Zubrycki and G. Granosik, "Novel haptic glove-based interface using jamming principle," in *Robot Motion and Control (RoMoCo), 2015 10th International Workshop on*. IEEE, 2015, pp. 46–51.

[18] T. M. Simon, R. T. Smith, and B. H. Thomas, "Wearable jamming mitten for virtual environment haptics," in *Proceedings of the 2014 ACM International Symposium on Wearable Computers*. ACM, 2014, pp. 67–70.

[19] M. R. Cutkosky, "On grasp choice, grasp models, and the design of hands for manufacturing tasks," *Robotics and Automation, IEEE Transactions on*, vol. 5, no. 3, pp. 269–279, 1989.

[20] T. Feix, J. Romero, H.-B. Schmiedmayer, A. M. Dollar, and D. Kragic, "The grasp taxonomy of human grasp types."

[21] K. Salisbury, D. Brock, T. Massie, N. Swarup, and C. Zilles, "Haptic rendering: Programming touch interaction with virtual objects," in *Proceedings of the 1995 symposium on Interactive 3D graphics*. ACM, 1995, pp. 123–130.

[22] D. C. Ruspini, K. Kolarov, and O. Khatib, "The haptic display of complex graphical environments," in *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*. ACM Press/Addison-Wesley Publishing Co., 1997, pp. 345–352.

[23] I. M. Vogels, "Detection of temporal delays in visual-haptic interfaces," *Human factors: The journal of the Human Factors and Ergonomics society*, vol. 46, no. 1, pp. 118–134, 2004.

[24] A. J. Doxon, D. E. Johnson, H. Z. Tan, and W. Provancher, "Human detection and discrimination of tactile repeatability, mechanical backlash, and temporal delay in a combined tactile-kinesthetic haptic display system," *Haptics, IEEE Transactions on*, vol. 6, no. 4, pp. 453–463, 2013.

[25] A. B. Swanson, I. B. Matev, and G. De Groot, "The strength of the hand," *Bull Prosthet Res*, vol. 10, no. 14, pp. 145–153, 1970.

[26] M. Turner, D. Gomez, M. Tremblay, and M. Cutkosky, "Preliminary tests of an arm-grounded haptic feedback device in telemanipulation," in *Proc. of the ASME Dynamic Systems and Control Division*, vol. 64, 1998, pp. 145–149.

[27] H. Z. Tan, X. D. Pang, and N. I. Durlach, "Manual resolution of length, force, and compliance," *Advances in Robotics*, vol. 42, pp. 13–18, 1992.

[28] C.-H. Ho and M. A. Srinivasan, "Human haptic discrimination of thickness," 1997.

[29] B. Calli, A. Walsman, A. Singh, S. Srinivasa, P. Abbeel, and A. M. Dollar, "Benchmarking in manipulation research: The ycb object and model set and benchmarking protocols," *arXiv preprint arXiv:1502.03143*, 2015.