High-Performance Wide-Area Optical Tracking
The HiBall Tracking System

Greg Welch, Gary Bishop, Leandra Vicci, Stephen Brumback, and Kurtis Keller:

University of North Carolina at Chapel Hill
Department of Computer Science, CB# 3175
Chapel Hill, NC 27599-3175 USA
01-919-962-1700
{welch, gb, vicci, brumback, keller}@cs.unc.edu

D'nardo Colucci:

Alternate Realities Corporation
27 Maple Place
Minneapolis, MN 55401 USA
01-612-616-9721
colucci@virtual-reality.com

Meta Platforms, Inc.
Exhibit 1103
Page 001

High-Performance Wide-Area Optical Tracking
The HiBall Tracking System

ABSTRACT

Since the early 1980's the Tracker Project at the University of North Carolina at Chapel Hill has been working on wide-area head tracking for Virtual and Augmented Environments. Our long-term goal has been to achieve the high performance required for accurate visual simulation throughout our entire laboratory, beyond into the hallways, and eventually even outdoors.

In this article we present results and a complete description of our most recent electro-optical system, the *HiBall Tracking System*. In particular we discuss motivation for the geometric configuration, and describe the novel optical, mechanical, electronic, and algorithmic aspects that enable unprecedented speed, resolution, accuracy, robustness, and flexibility.

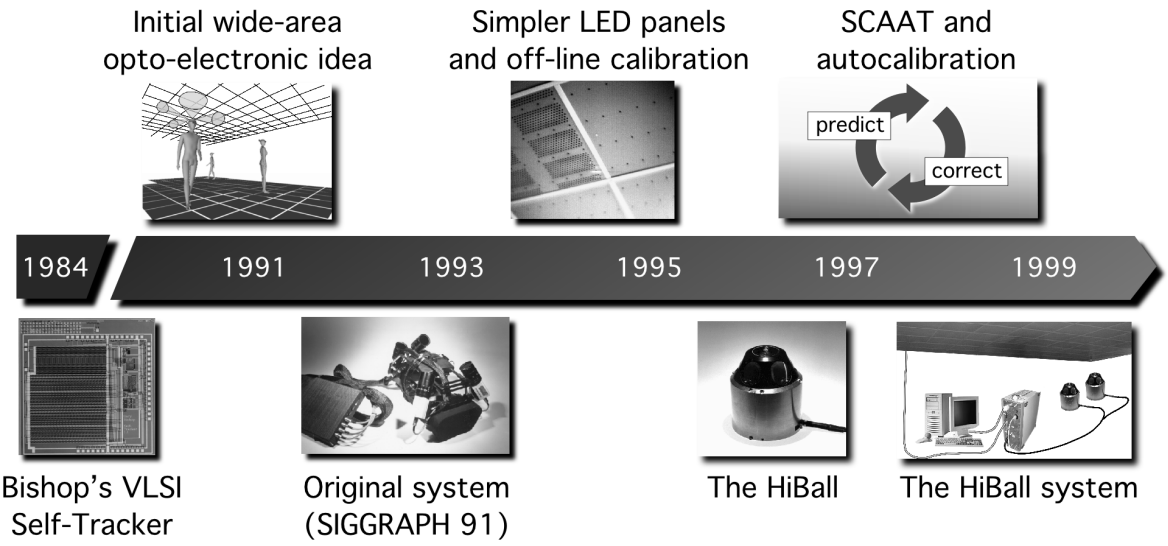

Figure 1

1. INTRODUCTION

Systems for *head tracking* for interactive computer graphics have been explored for over 30 years (Sutherland, 1968). As illustrated in Figure 1, the authors have been working on the problem for over twenty years (Azuma, 1993, 1995; Azuma & Bishop, 1994a, 1994b; Azuma & Ward, 1991; Bishop, 1984; Gottschalk & Hughes, 1993; UNC Tracker Project, 2000; Wang, 1990; J.-F. Wang et al., 1990; Ward, Azuma, Bennett, Gottschalk, & Fuchs, 1992; Welch, 1995, 1996; Welch & Bishop, 1997; Welch et al., 1999). From the beginning our efforts have been targeted at *wide-area* applications in particular. This focus was originally motivated by applications for which we believed that actually walking around the environment would be superior to virtually "flying." For example, we wanted to interact with room-filling virtual molecular models, and to naturally explore life-sized virtual architectural models. Today we believe that a wide-area system with high performance everywhere in our laboratory provides increased flexibility for all of our graphics, vision, and interaction research.

Meta Platforms, Inc.
Exhibit 1103
Page 002

1.1 Previous Work

In the early 1960's Ivan Sutherland implemented both mechanical and ultrasonic (carrier phase) head tracking systems as part of his pioneering work in virtual environments. He describes these systems in his seminal paper "A Head-Mounted Three Dimensional Display" (Sutherland, 1968). In the ensuing years, commercial and research teams have explored mechanical, magnetic, acoustic, inertial, and optical technologies. Complete surveys include (Bhatnagar, 1993; Burdea & Coiffet, 1994; Meyer, Applewhite, & Biocca, 1992; Mulder, 1994a, 1994b, 1998). Commercial magnetic tracking systems for example (Ascension, 2000; Polhemus, 2000) have enjoyed popularity as a result of a small user-worn component and relative ease of use. Recently inertial hybrid systems (Foxlin, Harrington, & Pfeifer, 1998; Intersense, 2000) have been gaining popularity for similar reasons, with the added benefit of reduced high-frequency noise and direct measurements of derivatives.

An early example of an optical system for tracking or motion capture is the *Twinkle Box* by Burton (Burton, 1973; Burton & Sutherland, 1974). This system measured the positions of user-worn flashing lights with optical sensors mounted in the environment behind rotating slotted disks. The *Selspot* system (Woltring, 1974) used fixed camera-like photo-diode sensors and target-mounted infrared light-emitting diodes that could be tracked in a one-cubic-meter volume. Beyond the HiBall Tracking System, examples of current optical tracking and motion capture systems include the *FlashPoint*© and *Pixsys*™ systems by Image Guided Technologies (IGT, 2000), the *laserBIRD*™ system by Ascension Technology (Ascension, 2000), and the *CODA Motion Capture System* by B & L Engineering (BL, 2000). These systems employ analog optical sensor systems to achieve relatively high sample rates for a moderate number of targets. Digital cameras (two-dimensional image-forming optical devices) are used in motion capture systems such as the *HiRes 3D Motion Capture System* by the Motion Analysis Corporation (Kadaba & Stine, 2000; MAC, 2000) to track a relatively large number of targets, albeit at a relatively low rate because of the need for 2D image processing.

1.2 Previous Work at UNC-Chapel Hill

As part of his 1984 dissertation on *Self-Tracker*, Bishop put forward the idea of outward looking tracking systems based on user-mounted sensors that estimate user *pose*[1] by observing landmarks in the environment (Bishop, 1984). He described two kinds of landmarks: high signal-to-noise-ratio beacons such as LEDs (light emitting diodes) and low signal-to-noise-ratio landmarks such as naturally occurring features. Bishop designed and demonstrated custom VLSI chips (Figure 2) that combined image sensing and processing on a single chip (Bishop & Fuchs, 1984). The idea was to combine multiple instances of these chips into an outward-looking cluster that estimated cluster motion by observing natural features in the un-modified environment. Integrating the resulting motion to estimate pose is prone to accumulating error, so further development required a complementary system based on easily detectable landmarks (LEDs) at known locations. This LED-based system was the subject of a 1990 dissertation by Jih-Fang Wang (Wang, 1990).

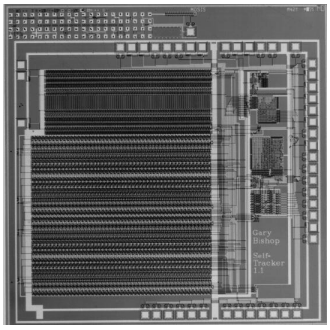

Figure 2

[1] We use the word *pose* to indicate both position and orientation (six degrees of freedom).

Meta Platforms, Inc.
Exhibit 1103
Page 003

Figure 3

In 1991 we demonstrated a working scalable electro-optical head-tracking system in the *Tomorrow's Realities* gallery at that year's ACM SIGGRAPH conference (J.-F. Wang et al., 1990; Wang, Chi, & Fuchs, 1990; Ward et al., 1992). The system (Figure 3) used four head-worn lateral effect photo-diodes that looked upward at a regular array of infrared LEDs installed in precisely machined ceiling panels. A user-worn backpack contained electronics that digitized and communicated the photo-coordinates of the sighted LEDs. Photogrammetric techniques were used to compute a user's head pose using the known LED positions and the corresponding measured photo-coordinates from each LEPD sensor (Azuma & Ward, 1991). The system was ground-breaking in that it was unaffected by ferromagnetic and conductive materials in the environment, and the working volume of the system was determined solely by the number of ceiling panels. (See Figure 3, top.)

1.3 The HiBall Tracking System

In this article we describe a new and vastly improved version of the 1991 system. We call the new system the *HiBall Tracking System*. Thanks to significant improvements in hardware and software this HiBall system offers unprecedented speed, resolution, accuracy, robustness, and flexibility. The bulky and heavy sensors and backpack of the previous system have been replaced by a small *HiBall* unit (Figure 4, bottom). In addition, the precisely machined LED ceiling panels of the previous system have been replaced by looser-tolerance panels that are relatively inexpensive to make and simple to install (Figure 4, top; Figure 10). Finally, we are using an unusual Kalman-filter-based algorithm that generates very accurate pose estimates at a high rate with low latency, and simultaneously self-calibrates the system.

Figure 4

As a result of these improvements the HiBall Tracking System can generate over 2000 pose estimates per second, with less than one millisecond of latency, better than 0.5 millimeters and 0.03 degrees of absolute error and noise, everywhere in a 4.5 by 8.5 meter room (with over two meters of height variation). The area can be expanded by adding more panels, or by using checkerboard configurations which spread panels over a larger area. The weight of the user-worn HiBall is about 300 grams, making it lighter than one optical sensor in the 1991 system. Multiple HiBall units can be daisy-chained together for head or hand tracking, pose-aware input devices, or precise 3D point digitization throughout the entire working volume.

2. DESIGN CONSIDERATIONS

In all of the optical systems we have developed (see Section 1.2) we have chosen what we call an *inside-looking-out* configuration, where the optical sensors are on the (moving) user and the *landmarks* (e.g., LEDs) are fixed in the laboratory. The corresponding *outside-looking-in*

Meta Platforms, Inc.
Exhibit 1103
Page 004

alternative would be to place the landmarks on the user, and to fix the optical sensors in the laboratory. (One can think about similar outside-in and inside-out distinctions for acoustic and magnetic technologies.) The two configurations are depicted in Figure 5.

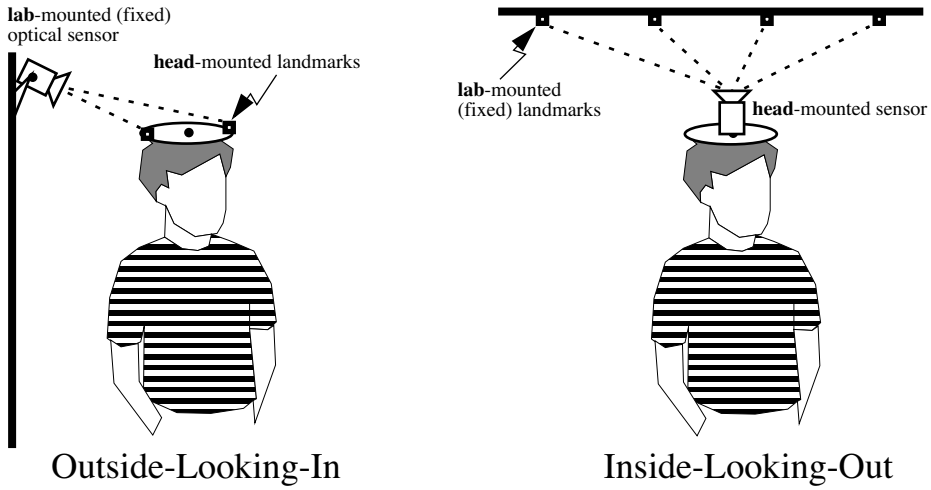

Figure 5

There are some disadvantages to the inside-looking-out approach. For small or medium-sized working volumes, mounting the sensors on the user is more challenging than mounting them in the environment. It is difficult to make user-worn sensor packaging small, and communication from the moving sensors to the rest of the system is more complex. In contrast, there are fewer mechanical considerations when mounting sensors in the environment for an *outside-looking-in* configuration. Because landmarks can be relatively simple, small, and cheap, they can often be located in numerous places on the user, and communication from the user to the rest of the system can be relatively simple or even unnecessary. This is particularly attractive for full-body motion capture (BL, 2000; MAC, 2000).

However there are some significant advantages to the inside-looking-out approach for head tracking. By operating with sensors on the user rather than in the environment, the system can be scaled indefinitely. The system can evolve from using dense active landmarks to fewer, lower signal-to-noise ratio, passive, and some day natural features for a Self-Tracker that operates entirely without landmark infrastructure (Bishop, 1984; Bishop & Fuchs, 1984; Welch, 1995).

The inside-looking-out configuration is also motivated by a desire to maximize sensitivity to changes in user pose. In particular, a significant problem with an outside-looking-in configuration is that only position estimates can be made directly, and so orientation must be inferred from position estimates of multiple fixed landmarks. The result is that orientation sensitivity is a function of both the *distance to the landmarks* from the sensor and the *baseline between the landmarks* on the user. In particular, as the distance to the user increases or the baseline between the landmarks decreases the sensitivity goes down. For sufficient orientation sensitivity one would likely need a baseline that is considerably larger than the user's head. This would be undesirable from an ergonomic standpoint and could actually restrict the user's motion.

With respect to translation, the change in measured photo-coordinates is the same for an environment-mounted (fixed) sensor and user-mounted (moving) landmark as it is for a user-mounted sensor and an environment-mounted landmark. In other words, the translation and corresponding sensitivity are the same for either case.

Meta Platforms, Inc.
Exhibit 1103
Page 005

3. SYSTEM OVERVIEW

The HiBall Tracking System consists of three main components (Figure 6). An outward-looking sensing unit we call the *HiBall* is fixed to each user to be tracked. The HiBall unit observes a subsystem of fixed-location infrared LEDs we call the *Ceiling*[1]. Communication and synchronization between the host computer and these subsystems is coordinated by the *Ceiling-HiBall Interface Board* (CIB). In Section 4 we describe these components in more detail.

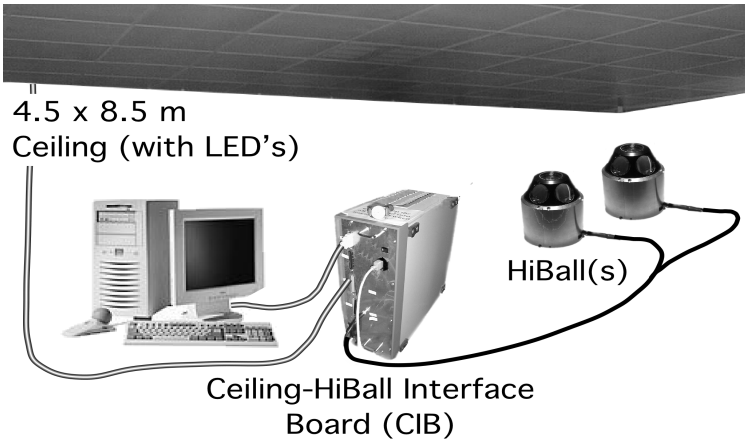

Figure 6

Each HiBall observes LEDs through multiple sensor-lens *views* that are distributed over a large solid angle. LEDs are sequentially flashed (one at a time) such that they are seen via a diverse set of views for each HiBall. Initial *acquisition* is performed using a brute force search through LED space, but once initial lock is made, the selection of LEDs to flash is tailored to the views of the active HiBall units. Pose estimates are maintained using a Kalman-filter-based prediction-correction approach known as *single-constraint-at-a-time* or SCAAT tracking. This technique has been extended to provide self-calibration of the Ceiling, concurrent with HiBall tracking. In Section 5 we describe the methods we employ, including the initial acquisition process and the SCAAT approach to pose estimation, with the *autocalibration* extension.

4. SYSTEM COMPONENTS

4.1 The HiBall

The original electro-optical tracker (Figure 3, bottom) used independently-housed lateral effect photo-diode units (LEPDs) attached to a light-weight tubular framework. As it turns out, the mechanical framework would flex (distort) during use, contributing to estimation errors. In part to address this problem the HiBall sensor unit was designed as a single rigid hollow ball having dodecahedral symmetry, with lenses in the upper six faces and LEPD on the insides of the opposing six lower faces (Figure 7). This immediately gives six primary "camera" *views* uniformly spaced by 57 degrees. The views efficiently share the same internal air space, and are rigid with respect to each other. In addition, light entering any lens sufficiently off axis can be

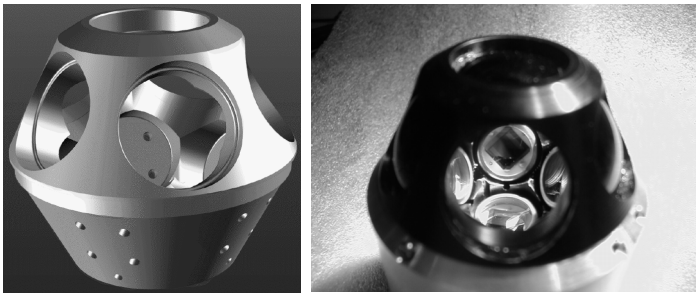

Figure 7

[1] At the present time, the LEDs are in fact entirely located in the ceiling of our laboratory, hence the sub-system name *Ceiling*, but LEDs could as well be located on walls or other fixed locations.

Meta Platforms, Inc.
Exhibit 1103
Page 006

seen by a neighboring LEPD, giving rise to five secondary views through the top or central lens, and three secondary views through the five other lenses. Overall, this provides 26 fields of view which are used to sense widely separated groups of LEDs in the environment. While the extra views complicate the initialization of the Kalman filter as described in Section 5.5, they turn out to be of great benefit during steady-state tracking by effectively increasing the overall HiBall field of view without sacrificing optical sensor resolution.

The lenses are simple plano-convex fixed-focus lenses. Infrared (IR) filtering is provided by fabricating the lenses themselves from RG-780 Schott glass filter material which is opaque to better than 0.001% for all visible wavelengths, and transmissive to better than 99% for IR wavelengths longer than 830 nm. The longwave filtering limit is provided by the DLS-4 LEPD silicon photodetector (UDT Sensors, Inc.) with peak responsivity at 950 nm but essentially blind above 1150 nm.

The LEPDs themselves are not imaging devices; rather they detect the centroid of the luminous flux incident on the detector. The x-position of the centroid determines the ratio of two output currents, while the y-position determines the ratio of two other output currents. The total output current of each pair are commensurate, and proportional to the total incident flux. Consequently, focus is not an issue, so the simple fixed-focus lenses work well over a range of LED distances from about half a meter to infinity. The LEPDs and associated electronic components are mounted on a custom rigid-flex printed circuit board (Figure 8). This arrangement makes efficient use of the internal HiBall volume while maintaining isolation between analog and digital circuitry, and increasing reliability by alleviating the need for inter-component mechanical connectors.

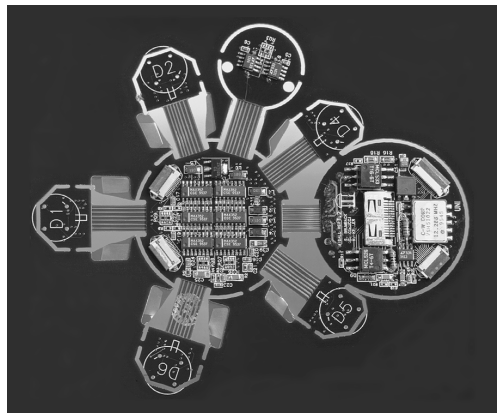

Figure 8

Figure 9 shows the physical arrangement of the folded electronics in the HiBall. Each LEPD has four transimpedance amplifiers (shown together as one "Amp" in Figure 9), the analog outputs of which are multiplexed with those of the other LEPDs, then sampled, held, and converted by four 16-bit Delta-Sigma analog-to-digital (A/D) converters. Multiple samples are integrated via an accumulator. The digitized LEPD data are organized into packets for communication back to the CIB. The packets also contain information to assist in error-detection. The communication protocol is simple, and while presently implemented by wire, the modulation scheme is amenable to a wireless implementation. The present wired implementation allows multiple HiBall units to be daisy-chained so a single cable can support a user with multiple HiBall units.

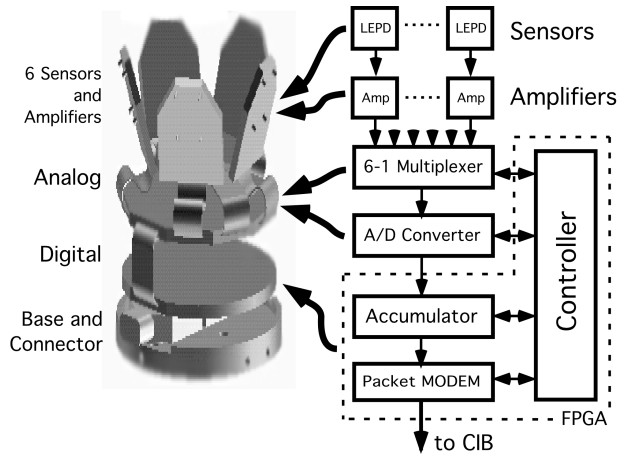

Figure 9

Meta Platforms, Inc.
Exhibit 1103
Page 007

4.2 The Ceiling

As presently implemented, the infrared LEDs are packaged in 61 centimeter square *panels*, to fit a standard false ceiling grid (Figure 10, top). Each panel uses five printed circuit boards: a main controller board and four identical transverse-mounted *strips* (bottom). Each strip is populated with eight LEDs for a total of 32 LEDs per panel. We mount the assembly on top of a metal panel such that the LEDs protrude through 32 corresponding holes. The design results in a Ceiling with a rectangular LED pattern with periods of 7.6 and 15.2 centimeters. This spacing is used for the initial estimates of the LED positions in the lab, then during normal operation the SCAAT algorithm continually refines the LED position estimates (Section 5.4). The SCAAT *autocalibration* not only relaxes design and installation constraints, but provides greater precision in the face of initial and ongoing uncertainty in the Ceiling structure.

Figure 10

We currently have enough panels to cover an area approximately 5.5 by 8.5 meters with a total of approximately 3,000 LEDs.[1] The panels are daisy-chained to each other, and panel selection encoding is position (rather than device) dependent. Operational commands are presented to the first panel of the daisy chain. At each panel, if the panel select code is zero the controller decodes and executes the operation; else it decrements the panel select code and passes it along to the next panel (controller). Upon decoding, a particular LED is selected and the LED is energized. The LED brightness (power) is selectable for *automatic gain control* as described in Section 5.2.

We currently use Siemens SFH-487P GaAs LEDs which provide both a wide angle radiation pattern and high peak power, emitting at a center wavelength of 880 nm in the near IR. These devices can be pulsed up to 2.0 Amps for a maximum duration of 200 μs with a 1:50 (on:off) duty cycle. While the current Ceiling architecture allows flashing of only one LED at a time, LEDs may be flashed in any sequence. As such no single LED can be flashed too long or too frequently. We include both hardware and software protection to prevent this.

4.3 The Ceiling-HiBall Interface Board

The Ceiling-HiBall Interface Board or CIB (Figure 11) provides communication and synchronization between a host personal computer, the HiBall (Section 4.1), and the Ceiling (Section 4.2). The CIB has four Ceiling ports allowing interleaving of ceiling panels for up to four simultaneous LED flashes and/or higher Ceiling bandwidth. (The Ceiling bandwidth is inherently limited by LED power restrictions as described in Section 4.2, but this can be increased by spatially multiplexing the Ceiling panels.) The CIB has two tether interfaces that can communicate with up to four daisy-chained HiBall units. The full-duplex communication with the HiBall units uses a modulation scheme (BPSK)

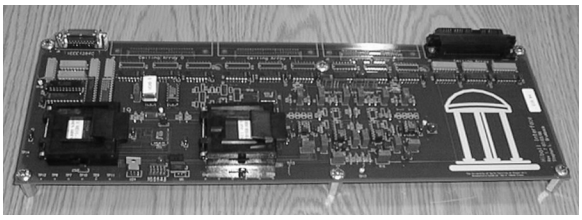

Figure 11

[1] The area is actually L-shaped; a small storage room occupies one corner.

Meta Platforms, Inc.
Exhibit 1103
Page 008

allowing future wireless operation. The interface from the CIB to the host PC is the stable IEEE1284C extended parallel port (EPP) standard.

The CIB comprises analog drive and receive components as well as digital logic components. The digital components implement store and forward in both directions and synchronize the timing of the LED "on" interval within the HiBall dark-light-dark intervals (Section 5.2). The protocol supports full-duplex flow control. The data are arranged into packets that incorporate error detection.

5. METHODS

5.1 Bench-Top (Off-Line) HiBall Calibration

After each HiBall is assembled we perform an off-line calibration procedure to determine the correspondence between image-plane coordinates and rays in space. This involves more than just determining the view transform for each of the 26 views. Non-linearities in the silicon sensor and distortions in the lens (e.g., spherical aberration) cause significant deviations from a simple pin-hole camera model. We dealt with all of these issues through the use of a two-part camera model. The first part is a standard pin-hole camera represented by a 3 x 4 matrix. The second part is a table mapping real image-plane coordinates to ideal image-plane coordinates.

Both parts of the camera model are determined using a calibration procedure that relies on a goniometer (an angular positioning system) of our own design. This device consists of two servo motors mounted together such that one motor provides rotation about the vertical axis while the second motor provides rotation about an axis orthogonal to vertical. An important characteristic of the goniometer is that the rotational axes of the two motors intersect at a point at the center of the HiBall optical sphere; this point is defined as the origin of the HiBall. (It is this origin that provides the reference for the HiBall state during run time as described in Section 5.3.) The rotational positioning motors were rated to provide 20 arc-second precision; we further calibrated them to 6 arc seconds using a laboratory grade theodolite—an angle measuring system.

In order to determine the mapping between sensor image-plane coordinates and three-space rays, we use a single LED mounted at a fixed location in the laboratory such that it is centered in the view directly out of the top lens of the HiBall. This ray defines the Z or up axis for the HiBall coordinate system. We sample other rays by rotating the goniometer motors under computer control. We sample each view with rays spaced about every 6 minutes of arc throughout the field of view. We repeat each measurement 100 times in order to reduce the effects of noise on the individual measurements and to estimate the standard deviation of the measurements.

Given the tables of approximately 2500 measurements for each of the 26 views, we first determine a 3 by 4 view matrix using standard linear least-squares techniques. Then we determine the deviation of each measured point from that predicted by the ideal linear model. These deviations are re-sampled into a 25 by 25 grid indexed by sensor-plane coordinates using a simple scan conversion procedure and averaging. Given a measurement from a sensor at run time (Section 5.2) we convert it to an "ideal" measurement by subtracting a deviation bilinearly interpolated from the nearest 4 entries in the table.

5.2 On-Line HiBall Measurements

Upon receiving a command from the CIB (Section 4.3), which is synchronized with a CIB command to the ceiling, the HiBall selects the specified LEPD and performs three measurements, one before the LED flashes, one during the LED flash, and one after the LED flash. Known as "dark-light-dark", this technique is used to subtract out DC bias, low frequency noise, and

Meta Platforms, Inc.
Exhibit 1103
Page 009

background light from the LED signal. We then convert the measured sensor coordinates to "ideal" coordinates using the calibration tables described in Section 5.1.

In addition, during run time we attempt to maximize the signal-to-noise ratio of the measurement with an automatic gain control scheme. For each LED we store a target signal strength factor. We compute the LED current and number of integrations (of successive accumulated A/D samples) by dividing this strength factor by the square of the distance to the LED, estimated from the current position estimate. After a reading we look at the strength of the actual measurement. If it is larger than expected we reduce the gain, if it is less than expected we increase the gain. The increase and decrease are implemented as on-line averages with scaling such that the gain factor decreases rapidly (to avoid overflow) and increases slowly. Finally we use the measured signal strength to estimate the noise on the signal using (Chi, 1995), and then use this as the measurement noise estimate for the Kalman filter (Section 5.3).

5.3 Recursive Pose Estimation (SCAAT)

The on-line measurements (Section 5.2) are used to estimate the pose of the HiBall during operation. The 1991 system collected a group of diverse measurements for a variety of LEDs and sensors, and then used a method of simultaneous non-linear equations called *Collinearity* (Azuma & Ward, 1991) to estimate the pose of the sensor fixture shown in Figure 3 (bottom). There was one equation for each measurement, expressing the constraint that a ray from the front principal point of the sensor lens to the LED, must be collinear with a ray from the rear principal point to the intersection with the sensor. Each estimate made use of a group of measurements (typically 20 or more) that together over-constrained the solution.

This *multiple constraint* method had several drawbacks. First, it had a significantly lower estimate rate due to the need to collect multiple measurements per estimate. Second, the system of non-linear equations did not account for the fact that the sensor fixture continued to move throughout the collection of the sequence of measurements. Instead the method effectively assumes that the measurements were taken simultaneously. The violation of this *simultaneity assumption* could introduce significant error during even moderate motion. Finally, the method provided no means to identify or handle unusually noisy individual measurements. Thus, a single erroneous measurement could cause an estimate to jump away from an otherwise smooth track.

In contrast, the approach we use with the new HiBall system produces tracker reports as each new measurement is made, rather than waiting to form a complete collection of observations. Because single measurements under-constrain the mathematical solution, we refer to the approach as *single-constraint-at-a-time* or SCAAT tracking (Welch, 1996; Welch & Bishop, 1997). The key is that the single measurements provide *some* information about the HiBall state, and thus can be used to incrementally improve a previous estimate. We intentionally fuse each individual "insufficient" measurement immediately as it is obtained. With this approach we are able to generate estimates more frequently, with less latency, with improved accuracy, and we are able to estimate the LED positions on-line concurrently while tracking the HiBall (Section 5.4).

We use a Kalman filter (Kalman, 1960) to fuse the measurements into an estimate of the HiBall *state* $\bar{x}$ (the pose of the HiBall). We use the Kalman filter—a minimum variance stochastic estimator—both because the sensor measurement noise and the typical user motion dynamics can be modeled as normally-distributed random processes, and because we want an efficient on-line method of estimation. A basic introduction to the Kalman filter can be found in Chapter 1 of (Maybeck, 1979), while a more complete introductory discussion can be found in (Sorenson, 1970), which also contains some interesting historical narrative. More extensive references can be

Meta Platforms, Inc.
Exhibit 1103
Page 010

found in (Brown & Hwang, 1992; Gelb, 1974; Jacobs, 1993; Lewis, 1986; Maybeck, 1979; Welch & Bishop, 1995). Finally, we maintain a Kalman filter web page (Welch & Bishop, 2000) with introductory, reference, and research material.

The Kalman filter has been used previously to address similar or related problems. See for example (Azarbayejani & Pentland, 1995; Azuma, 1995; Emura & Tachi, 1994; Fuchs (Foxlin), 1993; Mazuryk & Gervautz, 1995; Van Pabst & Krekel, 1993). A relevant example of a Kalman filter used for sensor fusion in wide-area tracking system is given by (Foxlin et al., 1998) which describes a hybrid inertial-acoustic system that is commercially-available today (Intersense, 2000).

The SCAAT approach is described in detail in (Welch, 1996; Welch & Bishop, 1997). Included there is discussion of the benefits of using the approach, as opposed to a *multiple-constraint* approach such as (Azuma & Ward, 1991). However one key benefit warrants discussion here. There is a direct relationship between the *complexity* of the estimation algorithm, the corresponding *speed* (execution time per estimation cycle), and the *change* in HiBall pose between estimation cycles (Figure 12). As the algorithmic complexity increases, the execution time increases, which allows for significant non-linear HiBall motion between estimation cycles, which in turn implies the need for a more complex estimation algorithm.

Figure 12

The SCAAT approach on the other hand is an attempt to reverse this cycle. Because we intentionally use a single constraint per estimate, the algorithmic complexity is drastically reduced, which reduces the execution time, and hence the amount of motion between estimation cycles. Because the amount of motion is limited we are able to use a simple dynamic (process) model in the Kalman filter, which further simplifies the computations. In short, the simplicity of the approach means it can run very fast, which means it can produce estimates very rapidly, with low noise.

The Kalman filter requires both a model of the process dynamics, and a model of the relationship between the process state and the available measurements. In part due to the simplicity of the SCAAT approach we are able to use a simple PV (position-velocity) process model (Brown & Hwang, 1992). Consider the simple example state vector $\bar{x}(t) = [x_p(t), x_v(t)]^T$, where the first element $x_p(t)$ is the pose (position or orientation) and the second element $x_v(t)$ is the corresponding velocity, i.e. $x_v(t) = \frac{d}{dt}x_p(t)$. We model the continuous change in the HiBall state with the simple differential equation

$$\frac{d}{dt}\bar{x}(t) = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} x_p(t) \\ x_v(t) \end{bmatrix} + \begin{bmatrix} 0 \\ \mu \end{bmatrix} u(t), \tag{1}$$

where $u(t)$ is a normally-distributed white (in the frequency spectrum) scalar noise process, and the scalar μ represents the magnitude or *spectral density* of the noise. We use a similar model with a distinct noise process for each of the six pose elements. We determine the individual noise magnitudes using an off-line simulation of the system and a non-linear optimization strategy that seeks to minimize the variance between the estimated pose and a known motion path. (See Section 6.2.2.) The differential equation (1) represents a continuous integrated random walk, or an integrated *Wiener* or *Brownian-motion* process. Specifically, we model each component of the

Meta Platforms, Inc.
Exhibit 1103
Page 011

linear and angular HiBall velocities as a random walk, and then use these (assuming constant inter-measurement velocity) to estimate the HiBall pose at time $t + \delta t$ as follows:

$$\bar{x}(t + \delta t) = \begin{bmatrix} 1 & \delta t \\ 0 & 1 \end{bmatrix} \bar{x}(t) \tag{2}$$

for each of the six pose elements. In addition to a relatively simple process model, the HiBall measurement model is relatively simple. For any Ceiling LED (Section 4.2) and HiBall view (Section 4.1), the 2D sensor measurement can be modeled as

$$\begin{bmatrix} u \\ v \end{bmatrix} = \begin{bmatrix} c_x / c_z \\ c_y / c_z \end{bmatrix} \tag{3}$$

where

$$\begin{bmatrix} c_x \\ c_y \\ c_z \end{bmatrix} = VR^{\mathrm{T}}(\dot{l}_{xyz} - \bar{x}_{xyz}), \tag{4}$$

V is the camera viewing matrix from Section 5.1, $\dot{l}_{xyz}$ is the position of the LED in the world, $\bar{x}_{xyz}$ is the position of the HiBall in the world, and R is a rotation matrix corresponding to the orientation of the HiBall in the world. In practice we maintain the orientation of the HiBall as a combination of a global (external to the state) quaternion and a set of incremental angles as described in (Welch, 1996; Welch & Bishop, 1997).

Because the measurement model (3)-(4) is non-linear we use an *extended Kalman filter*, making use of the Jacobian of the non-linear HiBall measurement model to transform the covariance of the Kalman filter. While this approach does not preserve the presumed Gaussian nature of the process, it has been used successfully in countless applications since the introduction of the (linear) Kalman filter. Based on observations of the statistics of the HiBall filter residuals, the approach also appears to work well for the HiBall. In fact it is reasonable to expect that it would, as the speed of the SCAAT approach minimizes the distance (in state space) over which we use the Jacobian-based linear approximation. This is another example of the importance of the relationship shown in Figure 12.

At each estimation cycle, the next of the 26 possible views is chosen randomly. Four points corresponding to the corners of the LEPD sensor associated with that view are projected into the world using the 3 by 4 viewing matrix for that view, along with the current estimates of the HiBall pose. This projection, which is the inverse of the measurement relationship described above, results in four rays extending from the sensor into the world. The intersection of these rays and the approximate plane of the Ceiling determines a 2D bounding box on the Ceiling, within which are the candidate LEDs for the current view. One of the candidate LEDs is then chosen in a least-recently-used fashion to ensure a diversity of constraints.

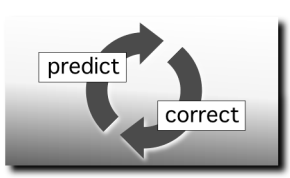

Once a particular view and LED have been chosen in this fashion, the CIB (Section 4.3) is instructed to flash the LED and take a measurement as described in Section 5.2. This single measurement is compared with a prediction obtained using (3), and the difference or *residual* is used to update the filter state and covariance matrices using the *Kalman gain* matrix. The Kalman gain is computed as a combination of the current

Meta Platforms, Inc.
Exhibit 1103
Page 012

filter covariance, the measurement noise variance (Section 6.2.1), and the Jacobian of the measurement model. This recursive prediction-correction cycle continues in an ongoing fashion, a single constraint at a time.

A more detailed discussion of the HiBall Kalman filter and the SCAAT approach is beyond the scope of this paper. For additional information see (Welch, 1996; Welch & Bishop, 1997).

5.4 On-line LED Autocalibration

Along with the benefit of simplicity and speed, the SCAAT approach offers the additional capability of being able to estimate the 3D positions of the LEDs in the world concurrently with the pose of the HiBall, on line, in real time. This capability is a tremendous benefit in terms of the accuracy and noise characteristics of the estimates. Accurate LED position estimates are so important that prior to the introduction of the SCAAT approach a specialized off-line approach was developed to address the problem (Gottschalk & Hughes, 1993).

The method we now use for autocalibration involves defining a distinct SCAAT Kalman filter for each LED. Specifically, for each LED we maintain a state $\hat{l}$ (estimate of the 3D position) and a 3x3 Kalman filter covariance. At the beginning of each estimation cycle we form an augmented state vector $\hat{x}$ using the appropriate LED state and the current HiBall state: $\hat{x} = [\bar{x}^T, \hat{l}^T]^T$. Similarly we augment the Kalman filter error covariance matrix with that of the LED filter. We then follow the normal steps outlined in Section 5.3, with the result being that the LED portion of the filter state and covariance is updated in accordance with the measurement residual. At the end of the cycle we extract the LED portions of the state and covariance from the augmented filter, and save them externally. The effect is that as the system is being used, it continually refines its estimates of the LED positions, thereby continually improving its estimates of the HiBall pose. Again, for additional information see (Welch, 1996; Welch & Bishop, 1997).

5.5 Initialization and Re-Acquisition

The recursive nature of the Kalman filter (Section 5.3) requires that the filter be initialized with a known state and corresponding covariance before steady-state operation can begin. Such an initialization or *acquisition* must take place prior to any tracking session, but also upon the (rare) occasion when the filter diverges and “loses lock” as a result of blocked sensor views for example.

The acquisition process is complicated by the fact that each LEPD sees a number of different widely separated views (Section 4.1). Therefore detecting an LED provides at best an ambiguous set of potential LED directions in HiBall coordinates. Moreover, before acquisition no assumptions can be made to limit the search space of visible LEDs. As such, a relatively slow brute-force algorithm is used to acquire lock.

We begin with an exhaustive LED scan of sufficiently fine granularity to ensure that the central primary field of view is not missed. For the present Ceiling, we flash every 13th LED in sequence, and look for it with the central LEPD until we get a hit. Then a sufficiently large patch of LEDs, centered on the hit, is sampled to ensure that several of the views of the central LEPD will be hit. The fields of view are disambiguated by using the initial hits to estimate the yaw of the HiBall (rotation about vertical), and finally more selective measurements are used to refine the acquisition estimate sufficiently to switch into tracking mode.

Meta Platforms, Inc.
Exhibit 1103
Page 013

6. RESULTS

Three days after the individual pieces of hardware were shown to be functioning properly we demonstrated a complete working system. After months of subsequent tuning and optimization, the system continues to perform both qualitatively and quantitatively as well, or in some respects *better*, than we had anticipated (Section 6.1). The articulation of this success is not meant to be self-congratulatory, but to give credit to the extensive and careful modeling and simulation performed prior to assembly (Section 6.2). In fact, the Kalman filter parameters found by the optimization procedure described in Section 6.2.2 were, and continue to be, used directly in the working system. Likewise much of the software written for the original simulations continues to be used in the working system.

6.1 On-Line Operation

The HiBall system is in daily use as a tool for education and research. For example, it was used by Martin Usoh et al. to perform Virtual Reality experiments comparing virtual "flying", walking in place, and real walking (Usoh et al., 1999). The researchers used the HiBall system to demonstrate that as a mode of locomotion, real walking is simpler, more straightforward, and more natural, than both virtual flying and walking in place. The unprecedented combination of large working volume and high performance of the HiBall system led the researchers to claim that there was nowhere else that they could have meaningfully performed the experiments.

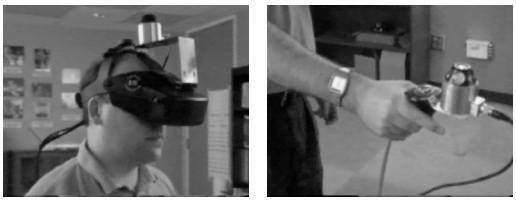

Figure 13

6.1.1 Robustness. As a result of a mechanical design trade-off, each sensor field of view is less than six degrees. The focal length is set by the size of the sensor housing, which is set by the diameter of the sensors themselves. Energetics is also a factor, limiting how small the lenses can be while maintaining sufficient light collecting area. As a result of these design trade-offs, even a momentary small error in the HiBall pose estimate can cause the recursive estimates to diverge and the system to lose lock after only a few LED sightings. And yet the system is quite robust. In practice users can jump around, crawl on the floor, lean over, even wave their hands in front of the sensors, and the system does not lose lock. During one session we were using the HiBall as a 3D *digitization probe*, a HiBall on the end of a pencil-shaped fiberglass wand (Figure 14, left). We laid the probe down on a table at one point, and were amazed to later notice that it was still tracking, even though it was only observing 3 or 4 LEDs near the edge of the Ceiling. We picked up the probe and continued using it, without it ever losing lock.

6.1.2 Estimate Noise. The simplest quantitative measurement of estimate noise is the standard deviation of the estimates when a HiBall is held stationary. With a tracker as sensitive as the HiBall it is important to be certain that it really is stationary. The raised floor in our laboratory allows motion, for example when a person walks by, that is larger than the expected error in the HiBall. We made careful measurements by resting the support for the HiBall on the concrete sub-floor in our laboratory. The standard deviation of the HiBall estimates while stationary was about 0.2 millimeters and 0.03 degrees. The distribution of the noise fit a normal distribution quite well.

To make measurements of the noise when the HiBall is in motion we rely on the assumption that almost all of the signal resulting from normal human motion is at frequencies below 2 Hz. We use a high-pass filter (Welch, 1967) on the pose estimates, and assume the output is noise. The resulting statistics are comparable to those made with the HiBall stationary, except at poses for

Meta Platforms, Inc.
Exhibit 1103
Page 014

which there are very few LEDs visible in only one or two views. In these poses, near the edge of the ceiling, the geometry of the constraints results in amplification of errors. For nearly all of the working volume of the tracker the standard-deviation of the noise on measurements while the HiBall is still or moving is about 0.2 millimeters and 0.03 degrees.

6.1.3 Absolute Accuracy. We have performed several experiments to measure the accuracy of the HiBall system, however the most objective experiment took place in July of 1999. Boeing Phantom Works scientists David Himmel and David Princehouse (Associate Technical Fellows) visited our laboratory for two days to assess the accuracy of the HiBall system, and its potential use in providing assembly workers with real-time feedback on the pose of hand-held pneumatic drills during the aircraft manufacturing process. (The right image in Figure 14 shows the HiBall attached to a pneumatic drill.)

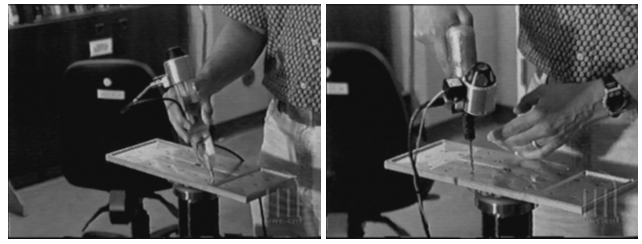

Figure 14

The scientists designed some controlled experiments to asses the accuracy of the HiBall system. They brought with them an aluminum "coupon" (see Figure 14 and Figure 15) with 27 shallow holes pre-drilled on 1.5 inch centers using a numerically-controlled milling machine with a stated accuracy of $1/1000$ inch. The holes (except one) were not actually drilled through the coupon, but instead formed conical dimples with a fine point at the center. The center-most hole (hole 14) was actually drilled completely through to provide a mounting point. Using that hole we attached the coupon to a military-grade tripod situated on the (false) floor of our laboratory, under the HiBall Ceiling. As shown in the left image of Figure 14 we mounted the HiBall on our standard probe, a rigid plastic pencil-like object with a pointed steel tip. We used one of the pre-drilled coupon holes to perform our normal HiBall probe calibration procedure, which involves placing the tip of the probe in the hole, pivoting the probe about the point while collecting several seconds of pose data, and then estimating the transformation from the HiBall to the probe tip. (We have a standard application that assists us with this procedure.) Together with Himmel and Princehouse we performed several experiments where we placed the tip of the HiBall probe in each hole in succession, sampling the HiBall pose estimates only when we pressed the probe button. We performed several such sessions over the course of one afternoon and the next morning (we re-calibrated the probe in the morning).

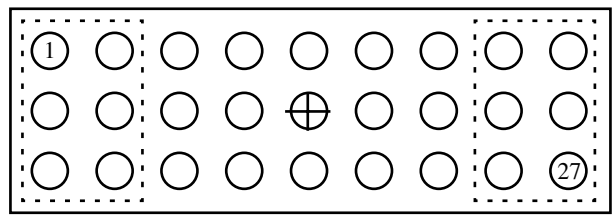

Figure 15

For the data from each session we used a least-squares optimization method to find an estimate of the full 6D transformation (translation and rotation) that minimized the Euclidian distance from the probe data to a 2D plane with 27 holes on 1.5 inch spacing. The resulting fit consistently corresponded to an average positioning error of $20/1000$ inch ($1/2$ millimeter) at the metal tip of the HiBall probe, which is within the target Boeing specifications. The system might actually be more accurate than our experiments indicated. For one, the diameter of the (rounded) tip of the HiBall probe is $1/2$ millimeter. In addition, at the time of the experiments we unfortunately did not heed our own advice to position the experimental platform on the rigid

Meta Platforms, Inc.
Exhibit 1103
Page 015

concrete sub-floor. In any case we are encouraged by the results, and are excited about the possibility that the HiBall system has uses beyond tracking for Virtual Reality.

6.2 Off-Line Simulation and Modeling

During the design of the HiBall system we made substantial use of simulation, in some domains to a very detailed level. For example, Zemax (Focus Software, 1995) was used extensively in the design and optimization of the optical design, including the design of the filter glass lenses, and geometry of the optical component layout. AutoCAD™ was used to design, specify, and fit-check the HiBall body mechanicals, to visualize the physical design, and to transmit the design to our collaborators at the University of Utah for fabrication by the *Alpha 1* System (Thomas, 1984; University of Utah Computer Science, 1999). A custom ray-tracing system was built by Stefan Gottschalk (UNC) for the purpose of evaluating the optical behavior and energetics of the primary, secondary, and tertiary fields of view; the results were used by the noise model developed in (Chi, 1995) as described in the next section.

In addition, a complete simulator of the system was written in C++. This simulator, discussed further in Section 6.2.2, was used to evaluate the speed, accuracy, and robustness of the system. In addition it was used to "tune" the Kalman filter for realistic motion dynamics. This simulator continues to be used to evaluate mechanical, optical, and algorithmic alternatives.

6.2.1 HiBall Measurement Noise Model. Signal-to-noise performance is a prime determiner of both accuracy and speed of the system, so an in-depth study (Chi, 1995) was performed to develop a detailed noise model accounting for properties of the LED, the LEPD (sensor), the optical system, the physical distance and pose, the electronics, and the dark-light-dark integrations described in Section 5.2. The predominant noise source is shot noise, with Johnson noise in the sheet resistivity of the LEPD surfaces being the next most significant. Careful measurements made in the laboratory with the actual devices yielded results that were almost identical to those predicted by the sophisticated model in (Chi, 1995). A simplified version of this model is used in the real system with the automatic gain control (Section 5.2) to predict the measurement noise for the Kalman filter (Section 5.3).

6.2.2 Complete System Simulations. To produce realistic data for developing and tuning our algorithms we collected several motion paths (sequences of pose estimates) from our first generation electro-optical tracker (Figure 3) at its 70 Hz maximum report rate. These paths were recorded from both naive users visiting our monthly "demo days" and from experienced users in our labs. In the same fashion as we had done for (Azuma & Bishop, 1994a) we filtered the raw path data with a non-causal zero-phase-shift low-pass filter to eliminate energy above 2 Hz. The output of the low-pass filtering was then re-sampled at whatever rate we wanted to run the simulated tracker, usually 1000 Hz. For the purposes of our simulations we considered these re-sampled paths to be the "truth"—a perfect representation of a user's motion. Tracking error was determined by comparing the "true" path to the estimated path produced by the tracker.

The simulator reads camera models describing the 26 views, the sensor noise parameters, the LED positions and their expected error, and the motion path described above. Before beginning the simulation, the LED positions are perturbed from their ideal positions by adding normally distributed error to each axis. Then, for each simulated cycle of operation, the "true" pose are updated using the input motion path. Next, a view is chosen and a visible LED within that view is selected, and the image-plane coordinates of the LED on the chosen sensor are computed using the camera model for the view and the LED as described in Section 5.3. These sensor coordinates are then perturbed based on the sensor noise model (Section 6.2.1) using the distance and angle to

Meta Platforms, Inc.
Exhibit 1103
Page 016

the LED. Now these noise corrupted sensor readings are fed to the SCAAT filter to produce an updated position estimate. The position estimate is compared to the true position to produce a scalar error metric described next.

The error metric we used combines the error in pose in a way that relates to the effects of tracker error on a head-worn display user. We define a set of points arrayed around the user in a fixed configuration. We compute two sets of coordinates for these points; the true position using the true pose, and their estimated position using the estimated pose. The error metric is then the sum of the distances between the true and estimated positions of these points. By adjusting the distance of the points from the user we can control the relative importance of the orientation and the position error in the combined error metric. If the distance is small, then the position error is weighted most heavily; if the distance is large then the orientation error is weighted most heavily. Our two error metrics for the entire run are the square-root of the sum of the squares of all the distances, and the peak distance.

6.2.3 Tuning. Determining the magnitudes of the SCAAT Kalman filter noise parameters (Section 5.3) is called *system identification* or *tuning*. We use Powell's method (Press, Teukolsky, Vetterling, & Flannery, 1990) to minimize the error metric described above. Starting with a set of parameters we run the simulator over a full motion run to determine the total error for the run. The optimizer makes a small adjustment to the parameters and the process is repeated. These runs required hours of computer time and some skill (and luck) in choosing the initial parameters and step sizes. Of course, it is important to choose motion paths that are representative of expected target motion. For example, a run in which the target is very still would result in very different tuning from a run in which the target moves very vigorously.

7. FUTURE WORK

7.1 Improving the HiBall

The current SCAAT filter form (Section 5.3) and tuning values (Section 6.2.3) are a compromise between the responsiveness desired for high dynamics, and the heavy filtering desired for smooth estimates during very slow or no motion. As such we are investigating the use of a multi-modal or *multiple-model* Kalman filter framework (Bar-Shalom & Li, 1993; Brown & Hwang, 1992). A multiple-model implementation of the HiBall should be able automatically, continuously, and smoothly choose between one Kalman filter tuned for high dynamics, and another tuned for little or no motion. We have this working in simulation, but not yet implemented in the real system.

As mentioned in Section 4.3, the system was designed to support wireless communication between the HiBall and the CIB, without significant modification or added information overhead. Despite the fact that commercial head-worn displays are themselves tethered at this time, we are beginning work on a completely wireless HiBall and head-worn display system. We also intend to use the wireless HiBall with projector-based displays where the user is otherwise wearing only polarized glasses. Furthermore the HiBall was designed with extra built-in digital input-output capabilities. We are considering possibilities for providing access to these signals for (wireless) user-centered input devices and even body-centric limb tracking.

Finally we note that a private startup company called 3rdTech (3rdTech, 2000) has negotiated a technology license with UNC for the existing HiBall Tracking System. 3rdTech is now marketing an updated system with simpler LED "strips" instead of ceiling panels.

Meta Platforms, Inc.
Exhibit 1103
Page 017

7.2 Wide-Field-of-View HiBall

Beyond improving the existing system, we continue to head down a path of research and development that will lead to systems with reduced dependency on the laboratory infrastructure. For example, our current Ceiling panel design with 32 LEDs per panel, provides far more dense coverage than we believe is necessary. The density of Ceiling LEDs is a result of design based on the original sensor fixture show in Figure 3. Given a more sparse field of LEDs we believe that we could achieve similar performance with a version of the HiBall that has a small number of *wide field of view* optical sensor units. This would further reduce the packaging size of the user-worn sensor component.

7.3 To the Hallway and Beyond

By leveraging the knowledge gained from successful work in the laboratory, our long term goal is to achieve similar performance with little or no explicit infrastructure, for example throughout a building or even (some day) outdoors. While high-performance 6D tracking outdoors is a tremendous challenge that is unlikely to be solved any time soon, we believe that the eventual solution will involve a clever and careful combination of multiple complementary technologies. In particular we are pursuing the hybrid approach initially presented in (Welch, 1995). We look forward to a day when high-performance 6D tracking outdoors enables pose-aware devices for work such as Feiner's outdoor augmented reality (Feiner, MacIntyre, Höllerer, & Webster, 1997; Höllerer, Feiner, Terauchi, Rashid, & Hallaway, 1999), the "WorldBoard" initiative (Spohrer, 1999a, 1999b), and other wonderful applications.

8. ACKNOWLEDGEMENTS

We acknowledge former Tracker Project members and contributors (alphabetically): Ronald Azuma, Henry Fuchs, Stefan Gottschalk, Pawan Kumar, John Thomas, Jih-Fang Wang, Mark Ward, Scott Williams, Mary Whitton, and Philip Winston. We thank Al Barr (California Institute of Technology) and John "Spike" Hughes (Brown University) for their contributions to the original off-line LED calibration work that led to the simpler Ceiling panels (Figure 1 and Figure 10). Finally we want to acknowledge our many collaborators in the NSF Science and Technology Center for Computer Graphics and Scientific Visualization (below), and in particular our collaborators in mechanical design and fabrication at the University of Utah: Rich Riesenfeld, Sam Drake, and Russ Fish.

This work was supported in part by DARPA/ETO contract number DABT 63-93-C-0048, "Enabling Technologies and Application Demonstrations for Synthetic Environments", Principal Investigators Frederick P. Brooks Jr. and Henry Fuchs (UNC); and by the National Science Foundation Cooperative Agreement no. ASC-8920219: "Science and Technology Center for Computer Graphics and Scientific Visualization," Center Director Rich Riesenfeld (University of Utah). Principal Investigators Al Barr (Caltech), Don Greenberg (Cornell University), Henry Fuchs (UNC), Rich Riesenfeld, and Andy van Dam (Brown University).

Meta Platforms, Inc.
Exhibit 1103
Page 018

REFERENCES

3rdTech. (2000, July 15). *3rdTech™* , [HTML]. 3rdTech. Available: http://www.3rdtech.com/ [2000, July 19].

Ascension. (2000). *Ascension Technology Corporation*, [HTML]. Ascension Technology Corporation. Available: http://www.ascension-tech.com/ [2000, September 15].

Azarbayejani, A., & Pentland, A. (1995). Recursive Estimation of Motion, Structure, and Focal Length. *IEEE Trans. Pattern Analysis and Machine Intelligence, 17*(6).

Azuma, R. T. (1993, July). Tracking Requirements for Augmented Reality. *Communications of the ACM, 36,* 50-51.

Azuma, R. T. (1995). *Predictive Tracking for Augmented Reality*. Unpublished Ph.D. Dissertation, University of North Carolina at Chapel Hill, Chapel Hill, NC USA.

Azuma, R. T., & Bishop, G. (1994a). A Frequency-Domain Analysis of Head-Motion Prediction, *Computer Graphics* (SIGGRAPH 94 Conference Proceedings ed., pp. 401-408). Los Angeles, CA: ACM Press, Addison-Wesley.

Azuma, R. T., & Bishop, G. (1994b). Improving Static and Dynamic Registration in an Optical See-Through HMD, *Computer Graphics* (SIGGRAPH 94 Conference Proceedings ed., pp. 197-204). Orlando, FL USA: ACM Press, Addison-Wesley.

Azuma, R. T., & Ward, M. (1991). *Space-Resection by Collinearity: Mathematics Behind the Optical Ceiling Head-Tracker* (Technical Report 91-048). Chapel Hill, NC USA: University of North Carolina at Chapel Hill.

Bar-Shalom, Y., & Li, X.-R. (1993). *Estimation and Tracking: Principles, Techniques, and Software*.: Artec House, Inc.

Bhatnagar, D. K. (1993). *Position trackers for Head Mounted Display systems: A survey* (Technical Report TR93-010). Chapel Hill, NC USA: University of North Carolina at Chapel Hill.

Bishop, G. (1984). *The Self-Tracker: A Smart Optical Sensor on Silicon*. Unpublished Ph.D. Dissertation, University of North Carlina at Chapel Hill, Chapel Hill, NC USA.

Bishop, G., & Fuchs, H. (1984, January 23-25). *The Self-Tracker: A Smart Optical Sensor on Silicon*. Paper presented at the Advanced Research in VLSI, Massachusetts Institute of Technology.

BL. (2000). *CODA mpx30 Motion Capture System*, [html]. B & L Engineering. Available: http://www.bleng.com/animation/coda/codamain.htm [2000, April 27].

Brown, R. G., & Hwang, P. Y. C. (1992). *Introduction to Random Signals and Applied Kalman Filtering* (Second ed.): Wiley & Sons, Inc.

Burdea, G., & Coiffet, P. (1994). *Virtual Reality Technology* (First ed.): John Wiley & Sons, Inc.

Burton, R. P. (1973). *Real-Time Measurement of Multiple Three-Dimensional Positions*., University of Utah, Salt Lake City, UT USA.

Burton, R. P., & Sutherland, I. E. (1974). Twinkle Box: Three-Dimensional Computer-Input Devices, *Proceedings of the National Computer Conference*. Chicago, IL USA.

Chi, V. L. (1995). *Noise Model and Performance Analysis Of Outward-looking Optical Trackers Using Lateral Effect Photo Diodes* (TR95-012). Chapel Hill, NC USA: University of North Carlina at Chapel Hill.

Emura, S., & Tachi, S. (1994). *Sensor Fusion based Measurement of Human Head Motion*. Paper presented at the 3rd IEEE International Workshop on Robot and Human Communication (RO-MAN 94 NAGOYA), Nagoya University, Nagoya, Japan.

Meta Platforms, Inc.
Exhibit 1103
Page 019

Feiner, S., MacIntyre, B., Höllerer, T., & Webster, A. (1997). A Touring Machine: Prototyping 3D Mobile Augmented Reality Systems for Exploring Urban Environments. *Personal Technologies, 1*(4), 208-217.

Focus Software. (1995). *ZEMAX Optical Design Program User's Guide, , Version 4.5*. Tucson, AZ USA.

Foxlin, E., Harrington, M., & Pfeifer, G. (1998). Constellation™ : A Wide-Range Wireless Motion-Tracking System for Augmented Reality and Virtual Set Applications. In M. F. Cohen (Ed.), *Computer Graphics* (SIGGRAPH 98 Conference Proceedings ed., pp. 371-378). Orlando, FL USA: ACM Press, Addison-Wesley.

Fuchs (Foxlin), E. (1993). *Inertial Head-Tracking (manual)*. Unpublished M.S. Thesis, Massachusetts Institute of Technology, Cambridge, MA USA.

Gelb, A. (1974). *Applied Optimal Estimation*. Cambridge, MA: MIT Press.

Gottschalk, S., & Hughes, J. F. (1993). Autocalibration for Virtual Environments Tracking Hardware, *Computer Graphics* (SIGGRAPH 93 Conference Proceedings ed.). Anaheim, CA USA: ACM Press, Addison Wesley.

Höllerer, T., Feiner, S., Terauchi, T., Rashid, G., & Hallaway, D. (1999). Exploring MARS: developing indoor and outdoor user interfaces to a mobile augmented reality system. *Computers & Graphics, 23*, 779-785.

IGT. (2000). *Image Guided Technologies*, [HTML]. Image Guided Technologies,. Available: http://www.imageguided.com/ [2000, September 15].

Intersense. (2000). *Intersense IS-900*, [html]. Intersense. Available: http://www.isense.com/ [2000, April 27].

Jacobs, O. L. R. (1993). *Introduction to Control Theory* (Second ed.): Oxford University Press.

Kadaba, M. P., & Stine, R. (2000). *Real-Time Movement Analysis Techniques and Concepts for the New Millennium in Sports Medicine*, [HTML]. Motion Analysis Corporation, Santa Rosa, CA USA. Available: http://www.motionanalysis.com/applications/movement/rtanalysis.html [2000, September 15].

Kalman, R. E. (1960). A New Approach to Linear Filtering and Prediction Problems. *Transaction of the ASME—Journal of Basic Engineering*, 35-45.

Lewis, F. L. (1986). *Optimal Estimation with an Introductory to Stochastic Control Theory.*: John Wiley & Sons, Inc.

MAC. (2000). *HiRes 3D Motion Capture System*, [html]. Motion Analysis Corporation. Available: http://www.motionanalysis.com/applications/movement/gait/3d.html [2000, September 15].

Maybeck, P. S. (1979). *Stochastic models, estimation, and control* (Vol. 141).

Mazuryk, T., & Gervautz, M. (1995). Two-Step Prediction and Image Deflection for Exact Head Tracking in Virtual Environments, *Proceedings of EUROGRAPHICS 95* (EUROGRAPHICS 95 ed., Vol. 14 (3), pp. 30-41).

Meyer, K., Applewhite, H., & Biocca, F. (1992). A Survey of Position Trackers. *Presence, a publication of the Center for Research in Journalism and Mass Communication*.

Mulder, A. (1994a). *Human Movement Tracking Technology* (Technical Report TR 94-1): School of Kinesiology, Simon Fraser University.

Mulder, A. (1994b, May 8, 1998). *Human Movement Tracking Technology: Resources*, [HTML]. School of Kinesiology, Simon Fraser University. Available: http://www.cs.sfu.ca/people/ResearchStaff/amulder/personal/vmi/HMTT.add.html [2000, September 15].

Meta Platforms, Inc.
Exhibit 1103
Page 020

Mulder, A. (1998, May 8, 1998). *Human Movement Tracking Technology*, [HTML]. School of Kinesiology, Simon Fraser University. Available: http://www.cs.sfu.ca/people/ResearchStaff/amulder/personal/vmi/HMTT.pub.html [2000, September 15].

Polhemus. (2000). *Polhemus*, [HTML]. Polhemus. Available: http://www.polhemus.com/home.htm [2000, September 15].

Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (1990). *Numerical Recipes in C; The Art of Scientific Computing* (Second ed.): Cambridge University Press.

Sorenson, H. W. (1970, July). Least-Squares estimation: from Gauss to Kalman. *IEEE Spectrum, 7,* 63-68.

Spohrer, J. (1999a). Information in Places. *IBM Systems Journal, Pervasive Computing, 38*(4).

Spohrer, J. (1999b, June 16). *WorldBoard; What Comes After the WWW?*, [HTML]. Learning Communities Group, ATG, (c)Apple Computer, Inc. Available: http://worldboard.org/pub/spohrer/wbconcept/default.html [1999, December 24, 1999].

Sutherland, I. E. (1968). A head-mounted three dimensional display, *Proceedings of the 1968 Fall Joint Computer Conference, AFIPS Conference Proceedings* (Vol. 33, part 1, pp. 757-764). Washington, D.C.: Thompson Books.

Thomas, S. W. (1984, December). *The Alpha_1 Computer-Aided Geometric Design System in the Unix Environment*. Paper presented at the Computer Graphics and Unix Workshop.

UNC Tracker Project. (2000, July 10). *Wide-Area Tracking; Navigation Technology for Head-Mounted Displays*, [HTML]. Available: http://www.cs.unc.edu/~tracker [2000, July 18].

University of Utah Computer Science. (1999). *Alpha 1 Publications*, [HTML]. University of Utah, Department of Computer Science. Available: http://www.cs.utah.edu/projects/alpha1/a1_publications.html [1999, May 28].

Usoh, M., Arthur, K., Whitton, M. C., Bastos, R., Steed, A., Slater, M., & Brooks, F. P., Jr. (1999). Walking > Walking-in-Place > Flying, in Virtual Environments. In A. Rockwood (Ed.), *Computer Graphics* (SIGGRAPH 99 Conference Proceedings ed., pp. 359-364). Los Angeles, CA USA: ACM Press, Addison Wesley.

Van Pabst, J. V. L., & Krekel, P. F. C. (1993, September 20 - 22). *Multi Sensor Data Fusion of Points, Line Segments and Surface Segments in 3D Space*. Paper presented at the 7th International Conference on Image Analysis and Processing—, Capitolo, Monopoli, Italy.

Wang, J.-F. (1990). *A real-time optical 6D tracker for head-mounted display systems*. Unpublished Ph.D. Dissertation, University of North Carolina at Chapel Hill, Chapel Hill, NC USA.

Wang, J.-F., Azuma, R. T., Bishop, G., Chi, V., Eyles, J., & Fuchs, H. (1990, April 16-20). *Tracking a Head-Mounted Display in a Room-Sized Environment with Head-Mounted Cameras*. Paper presented at the SPIE 1990 Technical Symposium on Optical Engineering and Photonics in Aerospace Sensing, Orlando, FL.

Wang, J.-f., Chi, V., & Fuchs, H. (1990, March 25-28). *A Real-time Optical 3D Tracker for Head-mounted Display Systems*. Paper presented at the Symposium on Interactive 3D Graphics, Snowbird, UT.

Ward, M., Azuma, R. T., Bennett, R., Gottschalk, S., & Fuchs, H. (1992, March 29 - April 1). *A Demonstrated Optical Tracker With Scalable Work Area for Head-Mounted Display Systems*. Paper presented at the Symposium on Interactive 3D Graphics, Cambridge, MA USA.

Welch, G. (1995). *Hybrid Self-Tracker: An Inertial/Optical Hybrid Three-Dimensional Tracking System* (TR95-048). Chapel Hill, NC, USA: University of North Carolina at Chapel Hill, Department of Computer Science.

Meta Platforms, Inc.
Exhibit 1103
Page 021

Welch, G. (1996). *SCAAT: Incremental Tracking with Incomplete Information*. Unpublished Ph.D. Dissertation, University of North Carolina at Chapel Hill, Chapel Hill, NC, USA.

Welch, G., & Bishop, G. (1995). *An Introduction to the Kalman Filter* (TR95-041). Chapel Hill, NC, USA: University of North Carolina at Chapel Hill, Department of Computer Science.

Welch, G., & Bishop, G. (1997). SCAAT: Incremental Tracking with Incomplete Information. In T. Whitted (Ed.), *Computer Graphics* (SIGGRAPH 97 Conference Proceedings ed., pp. 333-344). Los Angeles, CA, USA (August 3 - 8): ACM Press, Addison-Wesley.

Welch, G., & Bishop, G. (2000, January 23, 2000). *The Kalman Filter*, [html]. University of North Carlina at Chapel Hill. Available: http://www.cs.unc.edu/~welch/kalman/index.html [2000, April 29].

Welch, G., Bishop, G., Vicci, L., Brumback, S., Keller, K., & Colucci, D. n. (1999). The HiBall Tracker: High-Performance Wide-Area Tracking for Virtual and Augmented Environments, *Proceedings of the ACM Symposium on Virtual Reality Software and Technology* (pp. 1-11). University College London, London, United Kingdom (December 20 - 23): ACM SIGGRAPH, Addison-Wesley.

Welch, P. D. (1967). The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging Over Short, Modified Periodograms. *IEEE Transactions on Audio Electroacoust, AU*(15), 70-73.

Woltring, H. J. (1974). New Possibilities for Human Motion Studies by Real-Time Light Spot Position Measurement. *Biotelemetry, 1*, 132-146.

Meta Platforms, Inc.
Exhibit 1103
Page 022