

Declaration of Rachel J. Watters on Authentication of Publication

I, Rachel J. Watters, am a librarian, and the Head of Resource Sharing for the General Library System, Memorial Library, located at 728 State Street, Madison, Wisconsin, 53706. Part of my job responsibilities include oversight of Wisconsin TechSearch (“WTS”), an interlibrary loan department at the University of Wisconsin-Madison. I have worked as a librarian at the University of Wisconsin library system since 1998, starting as a graduate student employee in the Kurt F. Wendt Engineering Library and WTS, then as a librarian in Interlibrary Loan at Memorial Library. I began professional employment at WTS in 2002 and became WTS Director in 2011. In 2019, I became of Head of Resource Sharing for UW-Madison’s General Library System. I have a master’s degree in Library and Information Studies from the University of Wisconsin-Madison. Through the course of my studies and employment, I have become well informed about the operations of the University of Wisconsin library system, which follows standard library practices.

This Declaration relates to the dates of receipt and availability of the following:

Welch, G., Bishop, G., Vicci, L., Brumback, S., Keller, K., and Colucci, D. (February 2001). High-performance wide-area optical tracking: the HiBall Tracking System. *Presence: Teleoperators and Virtual Environments*, 10(1), 1-21.

Standard operating procedures for materials at the University of Wisconsin-Madison Libraries. When an issue was received by the Library, it would be checked in, stamped with the date of receipt, added to library holdings records, and made available

Declaration of Rachel J. Watters on Authentication of Publication

to readers as soon after its arrival as possible. The procedure normally took a few days or at most 2 to 3 weeks.

Exhibit A to this Declaration is a true and accurate copy of the title page with library date stamp of *Presence: Teleoperators and Virtual Environments* (February 2001), from the University of Wisconsin-Madison Library collection. Exhibit A also includes an excerpt of pages 1 to 21 of that issue, showing the article entitled *High-performance wide-area optical tracking: the HiBall Tracking System* (2001). Based on this information, the date stamp on the journal issue title page indicates *High-performance wide-area optical tracking: the HiBall Tracking System* (2001) was received by the Kurt F. Wendt Library, University of Wisconsin-Madison on May 28, 2001.

Based on the information in Exhibit A, it is clear that the journal issue was received by the library on or before May 28, 2001, catalogued and available to library patrons within a few days or at most 2 to 3 weeks after May 28, 2001.

Members of the interested public could locate the *Presence: Teleoperators and Virtual Environments* (February 2001) publication after it was cataloged by searching the public library catalog or requesting a search through WTS. The search could be done by title and/or subject key words. Members of the interested public could access the publication by locating it on the library's shelves or requesting it from WTS.

Declaration of Rachel J. Watters on Authentication of Publication

I declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code.

Date: December 13, 2021

Memorial Library
728 State Street
Madison, Wisconsin 53706



Rachel J. Watters
Head of Resource Sharing

Exhibit A

PRESENCE

TELEOPERATORS AND VIRTUAL ENVIRONMENTS

S
P9145

This resource is also available
on the WWW.
Use MADCAT to launch.



VOLUME 10 / NUMBER 1 / FEBRUARY 2001

MIT PRESS

META 1018
META V. THALES

PRESENCE
Teleoperators and
Virtual Environments

Volume 10, Number 1
February 2001
ISSN 1054-7460

Editor-in-Chief

Nathaniel I. Durlach
Virtual Environment and
Teleoperator Research Consortium
Research Laboratory of Electronics
MIT

Managing Editor

Rebecca Lee Gamett
MIT

Director of Corporate
Contributions

Thomas A. Furness III
Human Interface Technology
Laboratory
University of Washington

Editorial Address

Editor-in-Chief, *PRESENCE*
MIT
77 Massachusetts Avenue
Room 36-709
Cambridge, MA 02139
presence@mit.edu

Individuals wishing to submit
manuscripts should follow the
guidelines provided in this issue.

Abstracting and Indexing

PRESENCE: Teleoperators and Virtual Environments is included in *Computer Abstracts*, *Ergonomics Abstracts*, *Information Science Abstracts*, *International Aerospace Abstracts*, and *Multi-Index to Cyber-Space, Virtual, and Artificial Reality*. *CompuMath Citation Index*, *Current Contents/Engineering Computing and Technology*, *Research Alert*, and *SciSearch*.

Sponsorship

PRESENCE is indebted to the following organizations for their generous support:
Air Force Office of Scientific Research
Department of Biomedical Engineering, Boston University
Human Interface Technology Laboratory at the University of Washington
Office of Naval Research
Research Laboratory of Electronics, MIT
Visual Computing Department, Hewlett Packard Laboratories
Senior Editors
Woodrow Barfield
Virginia Polytechnic Institute
Gary Bishop
University of North Carolina

Larry Hodges
Georgia Institute of Technology
John M. Hollerbach
University of Utah
Randy Pausch
Carnegie Mellon University
Thomas B. Sheridan
MIT
Mel Slater
University College London
Kay Stanney
University of Central Florida
Elizabeth Wenzel
NASA Ames Research Center
David Zeltzer
Fraunhofer Center for Research in Computer Graphics
Michael Zyda
Naval Postgraduate School

Editorial Board

Bernard Adelstein
NASA Ames Research Center
Terry Allard
NASA Ames Research Center
Walter A. Aviles
Teneo Computing, LLC
Frank Biocca
Michigan State University
Jens Blauert
Ruhr-University Bochum,
Germany
Grigore Burdea
Rutgers University
H. Steven Colburn
Boston University
Rudy Darken
Naval Postgraduate School
Paul Dizio
Brandeis University
Stephen Ellis
NASA Ames Research Center
Scott S. Fisher
Telepresence Research
Richard M. Held
MIT
Kenneth O. Johnson
Johns Hopkins University
Lynette A. Jones
MIT
James R. Lackner
Brandeis University
Jaron Lanier
Columbia University
Susan Ledeman
Queen's University, Canada
Jack M. Loomis
University of California
Michael Macedonia
STRICOM
Michael Moshell
University of Central Florida
Michael Naimark
Interval Research Corporation
Albert "Skip" Rizzo
University of Southern California
Warren Robinett
University of North Carolina
Jannick Rolland
University of Central Florida
Roy Ruddle
University of Leeds

Karun Shimoga
CMU Robotics Institute
Barbara Shinn-Cunningham
Boston University
Gurminder Singh
National University of Singapore
Robert J. Stone
VR Solutions/Virtual Presence
Ltd. UK
Martin Stytz
Air Force Institute of
Technology
Susumu Tachi
University of Tokyo
James Templeman
Naval Research Laboratory
Geb Thomas
University of Iowa
Colin Ware
University of New Hampshire
Richard C. Waters
Mitsubishi Electric Research
Laboratory
Suzanne Weghorst
University of Washington
Janet Weisenberger
Ohio State University
Robert B. Welch
NASA Ames Research Center
Thomas E. von Wiegand
MIT

Business Offices and
Subscription Rates

PRESENCE: Teleoperators and Virtual Environments is published bimonthly (February, April, June, August, October, December) by The MIT Press, Cambridge, MA 02142-1407. Subscriptions and address changes should be addressed to MIT Press Journals, Five Cambridge Center, Cambridge, MA 02142-1407; (617) 253-2889; fax (617) 577-1545; e-mail journals-orders@mit.edu. An electronic, full-text version of *PRESENCE* is available from the MIT Press. Subscriptions are on a volume-year basis. Subscription rates:
Electronic only—Individuals \$45.00, Students/retired \$43.00, Institutions \$342.00. Canadians add the 7% GST. **Print and Electronic**—Individuals \$80.00, Students/retired \$48.00, Institutions \$380.00. Outside the U.S. and Canada add \$30.00 for postage and handling. Current issues are \$15.00. Back issue rates: Individuals \$32.00, Institutions \$64.00. Outside the U.S. and Canada add \$5.00 per issue for postage and handling. Canadians add 7% GST. Claims may be e-mailed to: journals-claims@mit.edu. Claims for missing issues will be honored free of charge if made within three months after the publication date of the issue. Prices subject to change without notice. <http://mitpress.mit.edu/PRES>

Postmaster

Send address changes to *PRESENCE*:

Teleoperators and Virtual Environments, Five Cambridge Center, Cambridge, MA 02142-1407. Periodicals postage paid at Boston, MA, and at additional post offices.

Advertising and Mailing List
Rental

Inquiries may be addressed to the Marketing Dept., MIT Press Journals, Five Cambridge Center, Cambridge, MA 02142-1407. (617) 253-2866; fax (617) 258-5028; e-mail: journals-info@mit.edu.

Copyright Information

Permission to photocopy articles for internal or personal use, or the internal or personal use of specific client, is granted by the copyright owner for users registered with the Copyright Clearance Center (CCC) Transactional Reporting Service, provided that the fee of \$8.00 per article-copy is paid directly to CCC, 222 Rosewood Drive, Danvers, MA 01923. The fee code for users of the Transactional Reporting Service is 1054-7460/01 \$8.00. For those organizations that have been granted a photocopy license with CCC, a separate system of payment has been arranged. Address all other inquiries to Subsidiary Rights Manager, MIT Press Journals, Five Cambridge Center, Cambridge, MA 02142-1407; e-mail: journals-rights@mit.edu.

PRESENCE: Teleoperators and Virtual Environments is available on microfilm from University Microfilms, Inc., 300 N. Zeeb Road, Ann Arbor, MI 48106.

© 2001 by the Massachusetts Institute of Technology.

PRESENCE

TELEOPERATORS AND VIRTUAL ENVIRONMENTS

VOLUME 10, NUMBER 1, FEBRUARY 2001

Editorial Notes	iii
Guest Editors' Introduction: VRST'99 Special Issue	iv

ARTICLES		
S	High-Performance Wide-Area Optical Tracking: The HiBall Tracking System	1
	<i>Greg Welch, Gary Bishop, Leandra Vicci, Stephen Brumback, Kurtis Keller, and D'nardo Colucci</i>	
S	GNU/MAVERIK: A Microkernel for Large-Scale Virtual Environments	22
	<i>Roger Hubbard, Jon Cook, Martin Keates, Simon Gibson, Toby Howard, Alan Murta, Adrian West, and Steve Pettifer</i>	
S	Patterns of Network and User Activity in an Inhabited Television Event	35
	<i>Chris Greenhalgh, Steve Benford, and Mike Craven</i>	
S	Components for Distributed Virtual Environments	51
	<i>Manuel Oliveira, Jon Crowcroft, and Mel Slater</i>	
S	An Adaptive Multiresolution Method for Progressive Model Transmission	62
	<i>Danny To, Rynson W. H. Lau, and Mark Green</i>	
S	Testbed Evaluation of Virtual Environment Interaction Techniques	75
	<i>Doug A. Bowman, Donald B. Johnson, and Larry F. Hodges</i>	
S	An Introduction to 3-D User Interface Design	96
	<i>Doug A. Bowman, Ernst Kruijff, Joseph J. LaViola, Jr., and Ivan Poupyrev</i>	
	An Overview of the COVEN Platform	109
	<i>Emmanuel Frécon, Gareth Smith, Anthony Steed, Mårten Stenius, and Olov Ståhl</i>	

FORUM

WHAT'S HAPPENING

128

UW-MADISON LIBRARY
COLLEGE OF ENGINEERING
MAY 2 - 2001
UW-MADISON, WI 5370

Greg Welch
welch@cs.unc.edu

Gary Bishop
gb@cs.unc.edu

Leandra Vicci
vicci@cs.unc.edu

Stephen Brumback
brumback@cs.unc.edu

Kurtis Keller
keller@cs.unc.edu
Department of Computer Science
University of North Carolina at
Chapel Hill

D'nardo Colucci
colucci@virtual-reality.com
Alternate Realities Corporation

High-Performance Wide-Area Optical Tracking

The HiBall Tracking System

Abstract

Since the early 1980s, 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.

I Introduction

Systems for head tracking for interactive computer graphics have been explored for more than thirty years (Sutherland, 1968). As illustrated in figure 1, the authors have been working on the problem for more than twenty years (Azuma, 1993, 1995; Azuma & Bishop, 1994a, 1994b; Azuma & Ward, 1991; Bishop, 1984; Gottschalk & Hughes, 1993; UNC Tracker Project, 2000; Wang, 1990; 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.

I.1 Previous Work

In the early 1960s, 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




Figure 1.

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); and 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, photodiode 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 *Flash-*

Point 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 2-D 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 pose¹ by observing landmarks in the environment (Bishop, 1984). He described two kinds of

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




Figure 2.

landmarks: high signal-to-noise-ratio beacons such as light-emitting diodes (LEDs) 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 unmodified 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).

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 (Wang et al., 1990; Wang, Chi, & Fuchs, 1990; Ward et al., 1992). The system (figure 3) used four, head-worn, lateral-effect photodiodes 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 ferromag-





Figure 3.

netic 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.)





Figure 4.

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 that simultaneously self-calibrates the system.

As a result of these improvements, the HiBall Tracking System can generate more than 2,000 pose estimates per second, with less than 1 ms of latency, better

than 0.5 mm and 0.03 deg. of absolute error and noise, everywhere in a 4.5 m \times 8.5 m room (with more than two meters of height variation). The area can be expanded by adding more panels, or by using checkerboard configurations that spread panels over a larger area. The weight of the user-worn HiBall is approximately 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 3-D 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, in which the optical sensors are on the (moving) user and the landmarks (for instance, the LEDs) are fixed in the laboratory. The corresponding *outside-looking-in* 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.

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




Figure 5.

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 explicit 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.

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*.² Communication and synchronization between the host computer and these subsystems is coordinated

2. At the present time, the LEDs are in fact entirely located in the ceiling of our laboratory (hence the subsystem name *Ceiling*), but LEDs could as well be located on walls or other fixed locations.




Figure 6.

by the *Ceiling-HiBall Interface Board (CIB)*. In section 4, we describe these components in more detail.

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 (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 photodiode units (LEPDs) attached to a lightweight 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 LEPDs on the insides of the opposing six lower faces (figure 7). This immediately gives six primary "camera" views uniformly spaced by 57 deg. 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 seen by a neighboring LEPD, giving rise to five secondary views through the top or central lens, and three secondary views





Figure 7.

through the five other lenses. Overall, this provides 26 fields of view that are used to sense widely separated groups of LEDs in the environment. Although 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 photo-detector (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, and the




Figure 8.

y-position determines the ratio of two other output currents. The total output current of each pair are commensurate and are 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 circuitboard (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 intercomponent mechanical connectors.

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.

4.2 The Ceiling

As presently implemented, the infrared LEDs are packaged in 61 cm 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 cm and 15.2 cm. 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.

We currently have enough panels to cover an area approximately 5.5 m by 8.5 m with a total of approximately 3,000 LEDs.³ 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

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




Figure 10.

controller decodes and executes the operation; otherwise, 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 μ s with a 1:50 (on:off) duty cycle. Although 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 (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




Figure 11.

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) 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 (Offline) 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. Nonlinearities in the silicon sensor and distortions in the lens (such as spherical aberration) cause significant deviations from a simple pinhole camera model. We dealt with all of these issues through the use of a two-part camera model. The first part is a standard pinhole camera represented by a

3×4 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 runtime as described in section 5.3.) The rotational positioning motors were rated to provide twenty arc-second precision; we further calibrated them to six arc seconds using a laboratory grade theodolite—an angle measuring system.

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 six minutes of arc throughout the field of view. We repeat each measurement 100 times to reduce the effects of noise on the individual measurements and to estimate the standard deviation of the measurements.

Given the tables of approximately 2,500 measurements for each of the 26 views, we first determine a 3×4 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 resampled into a 25×25 grid indexed by sensor-plane coordinates using a simple scan-conversion procedure and averaging. Given a measurement from a sensor at runtime (section 5.2), we convert it to an “ideal” measurement by subtracting a deviation bilinearly interpolated from the nearest four entries in the table.

5.2 Online 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 LED 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 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 runtime 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 online 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 online 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 nonlinear 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