

CHARACTERIZATION OF END-TO-END DELAYS IN HEAD-MOUNTED DISPLAY SYSTEMS

Mark R. Mine
University of North Carolina at Chapel Hill
3/23/93

1.0 INTRODUCTION

This technical report presents the results of measurements made to characterize the end-to-end delays in head-mounted display (HMD) systems. The term “end-to-end delay” is used here to describe the total time required for the displayed image in a HMD to change in response to the movement of a user’s head. Included in the total end-to-end delay is:

- 1) **Tracking system delay** - the time to measure the position and orientation of the user’s head.
- 2) **Application host delay** - execution time within the application host.
- 3) **Image generation delay** - time for the graphics engine to generate the resulting picture.
- 4) **Display system delay** - the time required to display the image in the head-mounted Display.

Primary motivation for this work was the desire to characterize the relative performance (in terms of measurement latency) of the various tracking technologies currently in use at the University of North Carolina in Chapel Hill (UNC-CH). Included in the test were four commercial magnetic trackers: two from Ascension Technology Corporation (the Ascension Bird and the newer Flock of Birds), and two from Polhemus Incorporated (the Polhemus 3Space Tracker and the 3Space Fastrak models). Also included in the test was UNC’s own Optical Ceiling Tracker. Note that both the Ascension Bird and the Polhemus 3 Space Tracker are older models and are no longer in production.

Though the main variable in all of the described measurements was the tracking technology in use, the utility of these measurements was not limited to the characterization of tracker measurement delays. These experiments provided an opportunity to determine the contribution of *all* the components in a HMD system to the end-to-end delays.

The experiments described in this paper are a part of the overall program at UNC to reduce the end-to-end delays in head-mounted display systems. It is believed that these delays are a key detractor from the feeling of presence in a virtual world. In addition, for see-through head-mounted displays, the presence of end-to-end delays greatly complicates the critical task of registration of virtual and real world objects. Reducing these system end-to-end delays, therefore, is one of the keys to improved head-mounted display system performance.

2.0 HMD SYSTEM END-TO-END DELAYS

Where do the end-to-end delays in a head-mounted display system come from?

For this paper, the end-to-end delay in an HMD system is defined as the time it takes for the displayed image in an HMD to change in response to a change in the position or orientation of a user’s head. The total end-to-end delay can be broken down into several major components, each of which represents a single stage in the “HMD pipeline”, the set of computations that take place in the generation of a virtual world. (see table 1 below).

Table 1: HMD Pipeline Delay Components

HMD Pipeline Stage	Delay Component
Spatial position/orientation sensing	Tracking System delay
Application host processing	Application Host delay
Graphics computation	Image Generation delay
Image display	Display System delay

These various stages of the HMD pipeline are depicted graphically in figure 1.

Spatial position/orientation sensing is the determination of the position and orientation of the user's head using one of the currently available tracking devices (includes mechanical, optical, acoustic, or magnetic trackers). This data is used in generating the user's view point relative to the virtual world.

Application host processing includes the collection of sensor data (such as, but not limited to, the tracking data above), and any other external input to the virtual world application. Application host processing also includes the execution of any simulation process that controls the behavior of objects within the virtual environment. Examples include the simulation of molecular interaction or the calculation of the inter body forces in an n-body simulation.

Graphics computation is the generation of the left and right eye images necessary for the presentation of a three dimensional image of the virtual world. In head-mounted display systems, the generated images depend upon both the information sensed via the tracking system (i.e. the position and orientation of the user's head) and any simulation information generated by the application process.

Image display is the presentation of the generated image to the user via the monitors embedded in the head-mounted display. Typically this will involve a raster scanout of the computed image data to LCD/CRTs mounted in front of the user's eyes.

Each stage contributes a portion of the total end-to-end delay in the system. The cost of communication between the various stages must also be included in calculating the total end-to-end delays.

Conceptually, the total end-to-end delay in an HMD system can be thought of as the time it takes one "packet" of tracking information to pass completely through the HMD pipeline. Beginning with the tracking system, a packet containing the most recently measured position and orientation of the user's head is sent to the application host. In the host, the "packet" is used to compute the updated viewpoint of the user relative to objects in the virtual world. Next, the "packet", in the form of the updated user viewpoint, is passed to the graphics engine and is used to generate the pixels of the image of the virtual world. Finally, the updated image is transmitted to the display system to be scanned out in the head-mounted display.

Currently, research in both the commercial and academic sectors into faster trackers, more powerful computers and high speed graphics engines is helping to reduce the contribution to the end-to-end delays of the first three stages of the HMD pipeline. Effort is currently underway at UNC in the development of techniques that compensate for the built in delay of typical raster display devices. (see Mine, Mark. and Bishop, Gary. Just-In-Time Pixels. UNC-Chapel Hill Department of Computer Science technical report TR93-005, January 1993)

END-TO-END DELAYS IN HEAD MOUNTED DISPLAY SYSTEMS

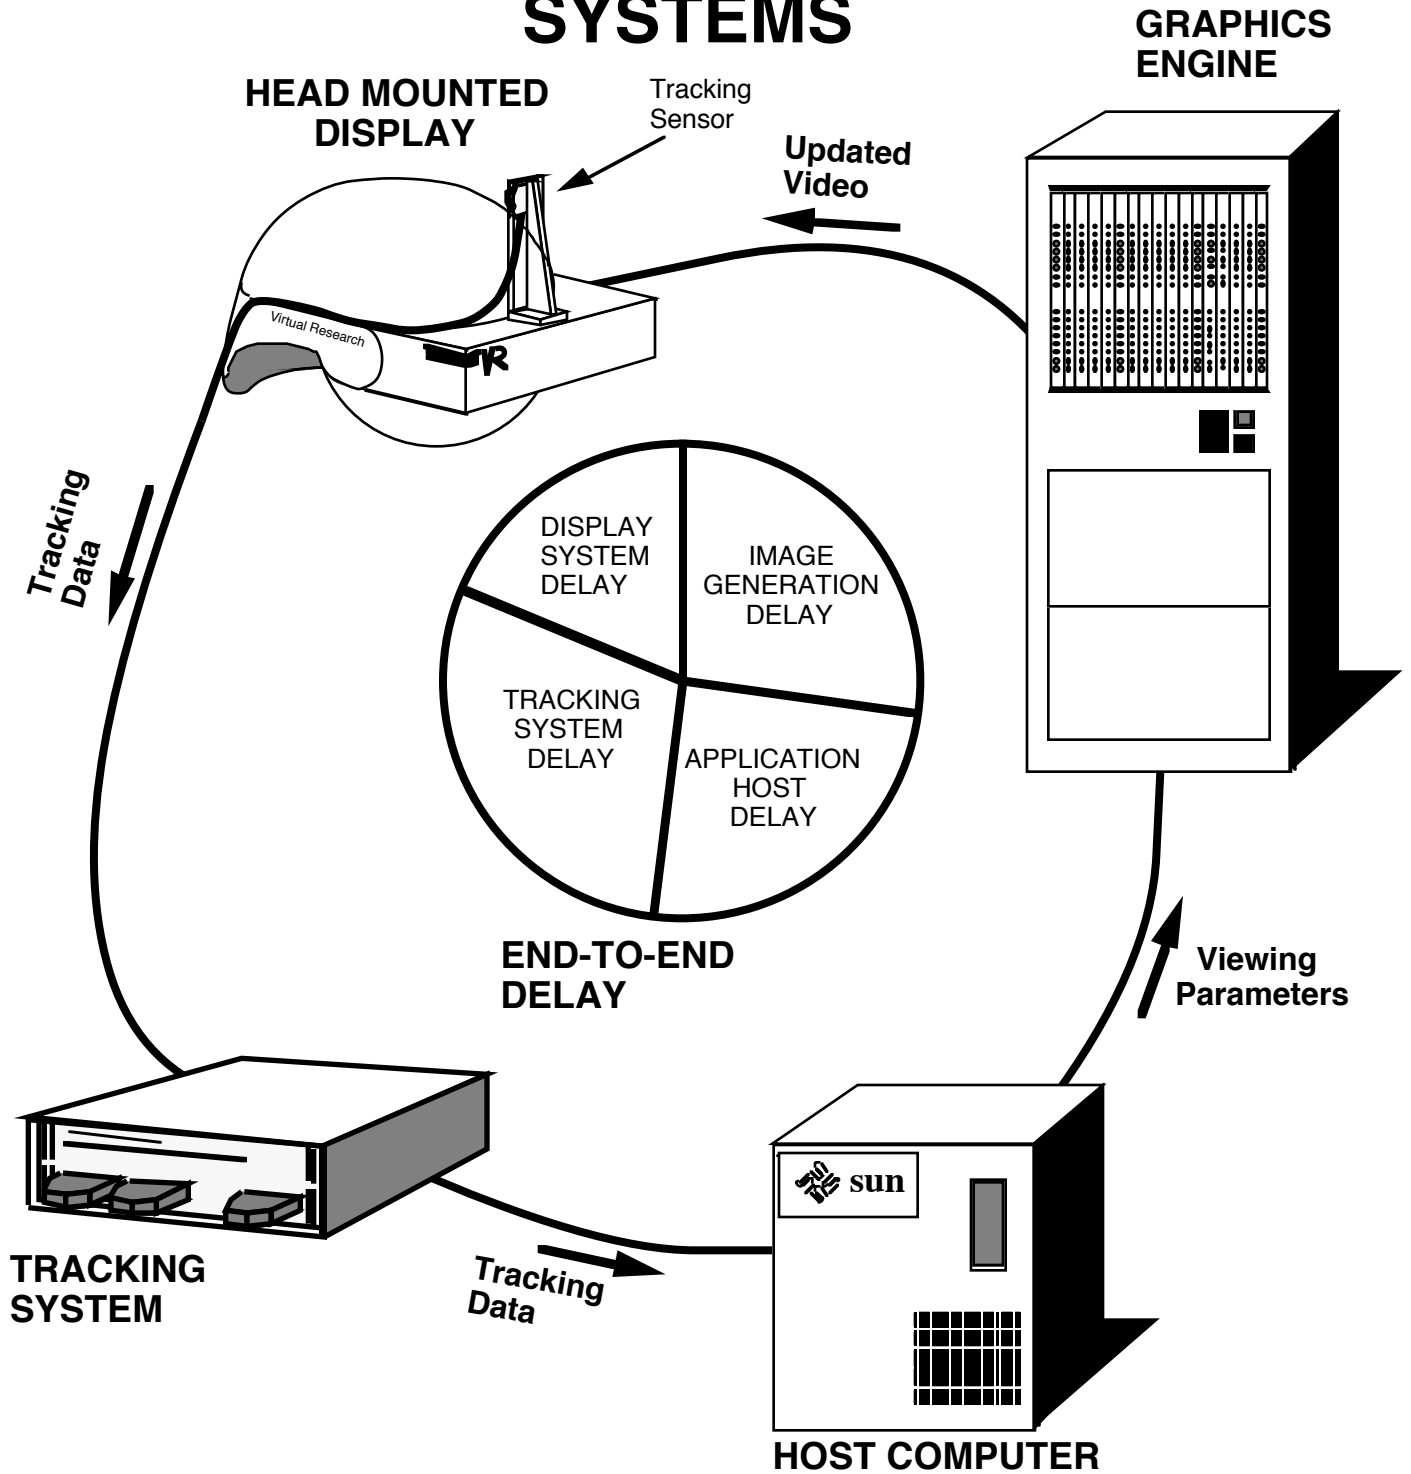


FIGURE 1

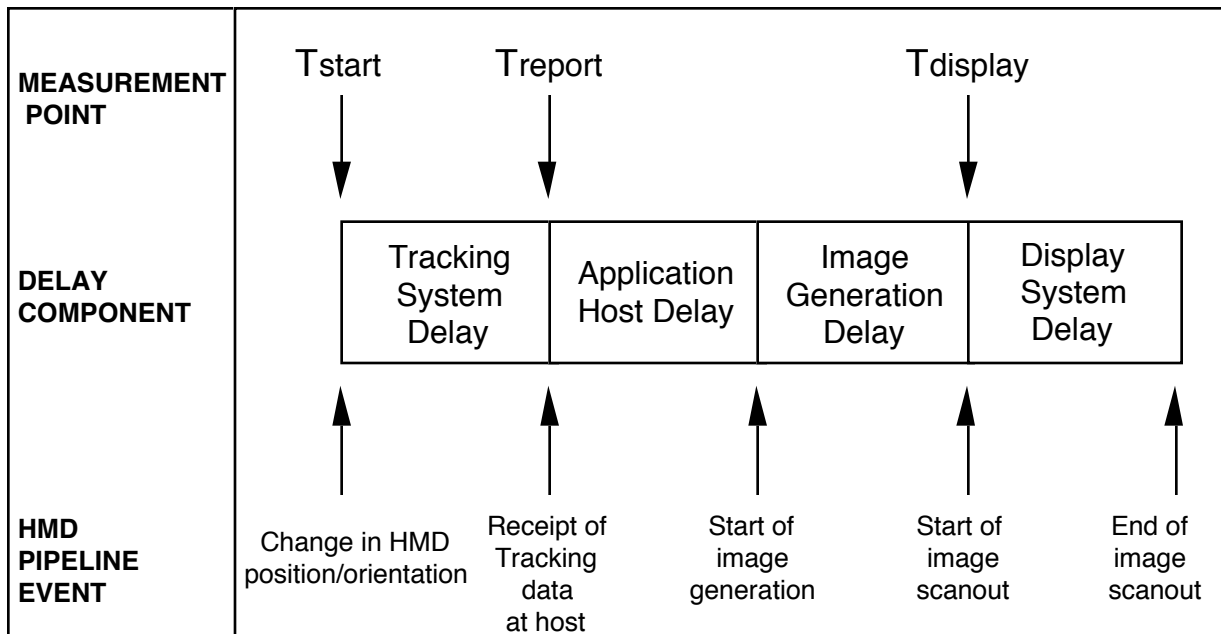
3.0 EXPERIMENTAL SETUP

To determine the end-to-end delays in an HMD system, measurement points were embedded in the HMD pipeline to monitor the transition of a single “packet” of tracking data from one stage of the pipeline to the next. Each measurement point was chosen to correspond to a critical event in the execution of the HMD pipeline:

- 1) Change in the position/orientation of the HMD
- 2) Receipt of tracking data at the host computer
- 3) Update of image in the HMD displays

As shown graphically in figure 2, these measurement points enable the determination of the relative contribution of each stage in the HMD pipeline to the total end-to-end delays of an HMD system. Note that no measurement was available to determine the start of image generation by the graphics engine. This means that in the resulting numbers, application host delay is combined with image generation delay. For the test application, however, application host delay was minimal (the only responsibility of the application host was the transfer of tracking data to the graphics engine - no simulation in support of the virtual world was required).

Figure 2: HMD Pipeline



Note: Relative size of each delay component not to scale

Figure 3 is a diagram of the experimental setup. Note that the tracking system sensor was actually removed from the HMD and mounted on the end of a pendulum arm. This was done to provide a known, repeatable pattern of motion which would facilitate the correlation between events in each stage of the pipeline. Also, a conventional monitor was used in place of the display screens mounted inside the HMD (to facilitate the placement of the test sensor).

Listed in table 2 are the various techniques that were used in detecting the transition of the tracking data “packet” from each stage in the HMD pipeline to the next:

END-TO-END DELAY MEASUREMENTS

FIGURE 3




Table 2: Delay Tests Measurement Techniques

Measurement Point	Label	Measurement Device
Change in HMD position/orientation	T_{start}	Photo diode/Pendulum
Receipt of tracking data at the host	T_{report}	Analog signal out host D/A port
Update of image in the HMD	T_{display}	Photo diode on Monitor

To aid in the correlation of “packets” in the various stages of the pipeline, a point in 3-dimensional space (which the tracking sensor passed through on each swing of the pendulum arm) was chosen as a point of reference. At each stage, the “packet” currently passing through the stage was compared with the reference point value and a signal was emitted when the packet corresponding to the point of reference (i.e. reporting the same position in 3-space) was received. Thus, detection of a signal at any of the measurement points listed in table 2 was indicative of the reference “packet” passing through that stage of the pipeline. The type of signal generated at each reference point and the means used for detecting that signal are discussed in more detail below.

Change in HMD position/orientation

The reference point chosen for use in all measurements was the low point on the arc of the pendulum’s motion. Mounted at this point was a photo diode/LED pair which was used to detect the physical transition of the pendulum arm through the reference point (by detecting the breaking of the beam of light between the photo diode/LED). By comparing the time the arm physically passed through the reference point (denoted T_{start} - see figure 2) with the time the tracking system reported that the arm had passed through that point (denoted T_{report}), a measure of the tracker latency was obtained. All signals used in the timing test were recorded using a digitizing oscilloscope (Tektronix DSA 602).

Receipt of Tracking Data at the Host

To determine when the “packet” corresponding to the reference point was received at the host (T_{report}), a record was made (at the start of each measurement cycle) of the position reported by the tracking system with the pendulum arm in the reference position. This recorded position was used by the host to compare with all subsequent tracker reports. For simplicity, the currently reported and reference positions were only compared along one axis (e.g. X, Y or Z). When the position reported by the tracking system passed through the reference point along the axis of measurement (i.e. changing from greater than to less than the reference value - or vice-versa) a signal was generated and sent to the oscilloscope for comparison with the signal generated by the pendulum.

The signal used to indicate the receipt of the reference tracking packet at the host (T_{report}) was generated using an Analog/Digital converter mounted in the host computer. The signal coming out of the A/D converter was modulated based upon the position currently reported by the tracking system. Two techniques were used, one merely indicating if the currently reported position is greater than or less than the reference position, the second indicating the current position reported by the tracking system (which results in a sine wave with the reference position at the zero point on the curve).

Update of Image in the HMD

To determine when the packet corresponding to the reference position was received at the displays, a single polygon was displayed on the screen and the color of the displayed

polygon was changed based upon the reported tracking position. When the currently reported position was greater than the reference position (along the axis of comparison) the polygon was colored white, when the reported tracking position was less than the reference position the polygon was colored black. The time of transition from black to white (or vice versa) was recorded using another photo diode mounted on the upper corner of the display (which was also connected to the oscilloscope) and is denoted T_{display} . The difference between T_{report} and T_{display} represents the time required to generate the image.

Image Scanout

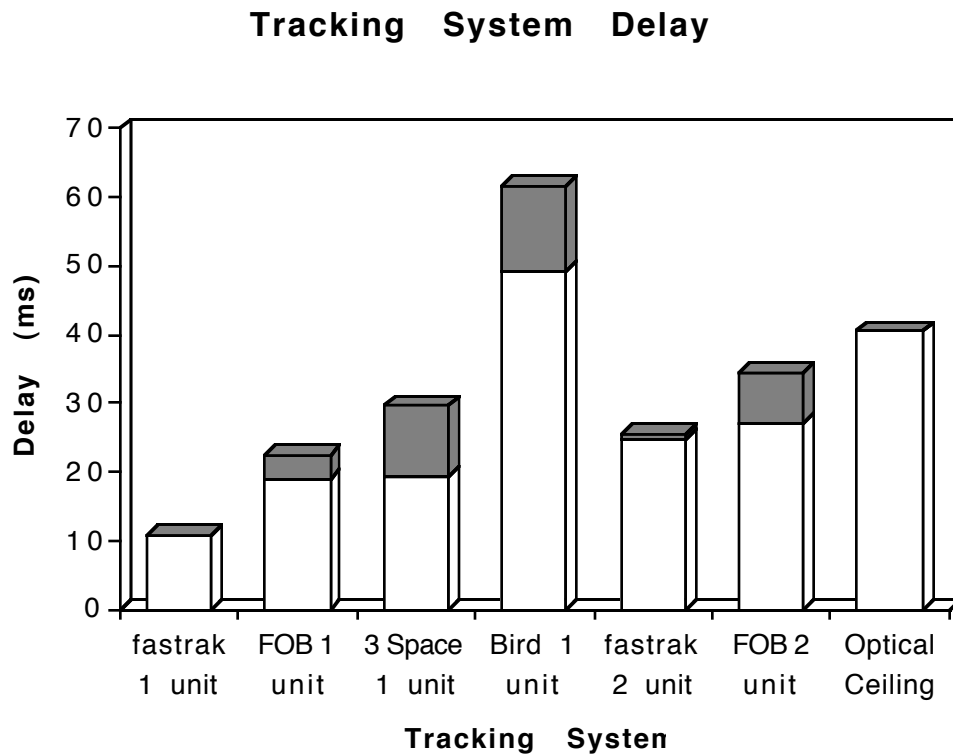
The time required to scanout an image depends upon the display device in use. For conventional NTSC based monitors, scanout of one field requires 16.67 ms.

4.0 EXPERIMENTAL RESULTS

4.1 Tracking System Delays

The first component of the end-to-end delay is the time required by the tracking system to measure the position and orientation of the user's head. This is the difference between T_{start} and T_{report} described above. Presented in graphically in figure 4 (and repeated in text form in table 3) is a summary of the average measured Tracking system delays.

Figure 4: Tracking System Measurement Delays



Notes:

- 1) Shaded area represents communication delays - see section 5.1
- 2) Flock of birds not in optimum communication configuration - see section 5.2.
- 3) Fastrak timing with position and orientation filter on - see section 5.3
- 4) Optical ceiling communication is via shared memory

In tabular form the results are:

Table 3: Tracking System Measurement Delays

TRACKING SYSTEM	DELAYS (MS)		
	<i>Tracker</i>	<i>Communication</i>	<i>Total</i>
Polhemus Fastrak - 1 Unit	10.65	0.3	10.95
Ascension Flock of Birds - 1 Unit	18.96	3.65	22.61
Polhemus 3Space - 1 Unit	19.23	10.4	29.63
Ascension Bird - 1 Unit	49.15	12.5	61.65
Polhemus Fastrak - 2 Units	24.9	0.6	25.50
Ascension Flock of Birds - 2 Units	27.09	7.3	34.39
UNC Optical Ceiling	40.5	-	40.5

Recall that the total delay (the difference between T_{start} and T_{report}) is the time between a change in the position/orientation of the HMD and the receipt of the corresponding tracking data at the host computer.

4.2 Application Host/Image Generation Delay

The next component of delay is the time between T_{report} and T_{display} . This includes the time spent executing the host program and the time required by the graphics engine to generate the image (including the time to copy the data to the frame buffer). In normal head-mounted display programs, the computation time of the host will include calculations and simulations required to determine the behavior of objects within the virtual world. Since the sole function of the host application for these timing tests was the comparison of incoming tracking data, this time was minimal.

For all of the experiments the same graphics engine was used (Pixel-Planes5), therefore a single average time to generate an image is applicable to all experiments.

The average Application Host/Image Generation time determined from these experiments was 54.36 ms. This is roughly equivalent to 3 NTSC field of latency. Note that this application host/image generation delay is for a very simple scene consisting of only one polygon. Greater delays can be expected for more complex scenes.

4.3 Display System Delay

The final component of the end-to-end delay is the contribution of the display system itself. All HMD systems currently in use at UNC make use of NTSC displays. In a standard NTSC system, images are generated at a rate of 30 Hz (i.e. one image every 33.34 ms). Since NTSC images are interlaced, (odd scanlines of an image being shown in one field followed by the even scanlines in the next field) the minimum display-system delay is 16.67 ms.

The display system delay is also dependent upon the type of display device in use. HMDs typically make use of CRT or LCD displays. See section 5.5 for further discussion of this topic.

4.4 Total Non-Tracking System Delay

The total non-tracking system delay (application host delay + graphics system delay + display system delay) is therefore approximately 70 ms from the receipt of tracking data at the host to the end of scanout of the first field of data.

5.0 ANALYSIS OF RESULTS

There are several factors to consider in analyzing the results:

5.1 Tracker Communication Delays

It is important to remember that the tracker delays reported in section 4.0 include both the communication time between the tracker and the host and the time required by the host to retrieve the data from internal buffers (see section 5.4 below). Communication time depends upon both the length of the tracker report sent from the tracker to the host and the speed of the connection between the two.

Table 4 lists the report record size, the communication link speed and the estimated communication time for the test configuration:

Table 4: Tracking System Report Size

Tracker	Report Size	Communication Link Speed	Communication Delays
Ascension Bird	240 bits/unit	19.2 K	12.5 ms
Ascension Flock	140 bits/unit	38.4 K	3.65 ms
Polhemus 3Space	200 bits/unit	19.2 K	10.4 ms
Polhemus Fastrak	31 bytes/unit	100Kbytes max	0.3 ms

Notes:

- 1) The reduction of the Ascension report size from 240 bits to 140 bits reflects the switch from position matrix report formats (12 word apiece) to position/quaternion formats (7 words apiece).
- 2) These communication delays do not include the time within the host to retrieve the data from internal buffers. This is only the time to physically transmit the data from the tracker to the host.

5.2 Flock of Birds Configuration

It should be noted that the best possible communication configuration was not used for the Flock of Birds in obtaining the results given in section 4.0. Only a single RS-232 line running at 38.4K was used to communicate with the flock for both commanding and receipt of data. This meant that in the two-unit configuration, reports from both sensors had to pass through the single RS-232 line. Since a separate electronics box is provided for each receiver, improved performance can be obtained by using a dedicated RS-232 line for each unit. Alternately, improved performance can be obtained by using a single dedicated RS485 connection running at 250K to talk to all birds in the flock. Due to the unavailability of such an interface on our host computer, this configuration was not tested.

5.3 Tracking System Filtering Parameters

Each tracker filters the incoming data before transmitting it to the host. Though this reduces the noise in the reported data it also significantly increases the tracker delays.

Since, for all the trackers tested, it is possible to specify the filtering level via software commands, it was decided that each tracker should be configured to use the minimum acceptable level of filtration. Though this served to minimize the tracking system delays, it is important to realize that in actual operation the level of filtration required may actually be higher than what was used in performing these tests. The use of these minimal filtering configurations will result in minimal tracking system delays but they may also result in unacceptably noisy data.

Table 5 list the tracker filtering parameters used in this test:

Table 5: Tracking System Filtering Parameters

Tracking System	Filter Parameters
Ascension Bird	DC Filter ON/AC NARROW
Ascension Flock of Birds	DC Filter ON/AC NARROW
Polhemus 3Space	Position/Orientation Filtering Off
Polhemus Fastrak	Position/Orientation Filtering On

5.4 Unix Interaction

One of the discoveries made in the course of these experiments is the effect of Unix buffering on the end-to-end delay results. Earlier measurements were affected by the fact that Unix would not process the data at the serial port as it arrived but would wait approximately 30 ms before servicing the port. Modifications had to be made to the kernel to ensure that the data was processed as it was received (i.e. an interrupt was received after each character instead of once every 30 ms). When the experiments were repeated with the kernel modification in place, a 20-25ms improvement was realized in the resulting end-to-end delays. All measurements results presented in this paper reflect the use of the modified Unix kernel (i.e. interrupt after every character received from the tracker).

5.5 Display System Peculiarities

The LCD displays mounted inside our current head-mounted displays were not used in gathering the data (intervening optics made it difficult to mount a photodiode directly on the screen). A conventional CRT monitor receiving the same video feed was used in its place. This made it possible to mount the photodiode directly on the screen and in the upper left hand corner of the display. This position allowed the detection of the start of the raster scan.

It is important to note, however, that the delay characteristics of LCDs that are used in most head-mounted displays are different from CRTs. Though they use the same NTSC video feed as the CRT monitors, rough measurements made here at UNC show that LCD monitors add at least one field time of delay (16.67 ms) over conventional CRTs.

6.0 CONCLUSIONS

Currently, most graphic systems are optimized for high throughput instead of minimum end-to-end delay. Though high throughput is obviously a desirable feature for head-mounted

display systems, the development of low-latency graphic systems must also be an important research goal. The presence of large delays between the start of head motion and the corresponding update of the displayed image greatly detracts from the feeling of presence in a virtual world. Furthermore the reduction of end-to-end delays is especially important in light of the growing interest in the development of see-through head-mounted displays. The presence of large end-to-end delays greatly complicates the task of registration of virtual and real world objects. The reduction of end-to-end delays, therefore, is a key step towards improved head-mounted display system performance.