Saccade Contingent Updating in Virtual Reality

Jochen Triesch
Cognitive Science Dept.
UC San Diego

Brian T. Sullivan
Center for Visual Science
U. of Rochester

Mary M. Hayhoe
Center for Visual Science
U. of Rochester

Dana H. Ballard
Computer Science Dept.
U. of Rochester

Abstract

We are interested in saccade contingent scene updates where the visual information presented in a display is altered while a saccadic eye movement of an unconstrained, freely moving observer is in progress. Since saccades typically last only several tens of milliseconds depending on their size, this poses difficult constraints on the latency of detection. We have integrated two complementary eye trackers in a virtual reality helmet to simultaneously 1) detect saccade onsets with very low latency and 2) track the gaze with high precision albeit higher latency. In a series of experiments we demonstrate the system s capability of detecting saccade onsets with sufficiently low latency to make scene changes while a saccade is still progressing. While the method was developed to facilitate studies of human visual perception and attention, it may find interesting applications in human-computer interaction and computer graphics.

CR Categories: H.5.2 [Information Interfaces and Presentation]: User Interfaces Input devices and Strategies; H.1.2 [Models and Principles]: User/Machine Systems Human information processing; I.3.6 [Computer Graphics]: Methodology and Techniques Interaction techniques

Keywords: eye tracking, saccades, saccade contingent updating, virtual reality, change blindness, limbus tracking

1 Introduction

Why saccade contingent scene updates? Saccade contingent updating (SCU), i.e. changing a presented visual stimulus while the eye is in motion, is an experimental technique that has proven quite valuable in studies of visual perception, in particular at addressing the nature of visual representations and memory, high level visual attention, and change blindness. While the eye is moving, image blur makes the visual system effectively blind. Saccadic suppression imposes an additional sensitivity loss for a period that outlasts the saccade by 50 ms or more [Stevenson et al. 1986]. Scientists can take advantage of these natural phenomena to examine how individuals respond to changes that occur in a visual scene while visual feedback is momentarily offline. This can be done without the concern that the subjects are merely detecting the change with low level sensory mechanisms. Our lab is particularly interested in the study of visual attention through the use of change blindness paradigms. A general overview of this approach could be described as follows: A visual stimulus is presented to a subject and during that period a significant change occurs in the visual scene. The subjects are typically informed to be on the lookout for changes, or are sometimes not informed but still questioned at the end of the experiment if they have noticed any visual changes occurring. A review of change blindness studies can be found in [Simons 2000]. To prevent low level mechanisms from immediately detecting this change in the visual presentation a transient is used to mask the change. To this end, previous change blindness experiments have used mud splashes, blank screen flashes, creative editing of films, and saccade contingent updating. It has been repeatedly demonstrated in these experiments that subjects will often be surprisingly unaware of even large changes occurring in the scene. These studies now play an important role in the debate on the nature of visual representations by indicating that humans seem to use rather minimal representations of incoming visual stimuli, or at least that consciousness has access only to a small part of this information. Unfortunately, many of these studies use rather unnatural experimental conditions making it difficult to generalize their results to everyday perception. Making changes during mud splashes or via video editing is somewhat unnatural and may not provide conclusive evidence about visual representations during natural, purposeful behavior in a 3-D world. Saccade contingent updating (SCU) on the other hand exploits a natural detail of human visual function so that changes in the visual scene can be made without the need to resort to unnatural interventions. Unfortunately, previous SCU setups have had to use bite bars or other devices to keep the subject s head fixed because the eye-tracking devices used were either too bulky to be portable and/or too sensitive to body movements. Because of this, visual presentations in these experiments are typically done on a computer monitor in 2-D. Our goal was to develop an experimental setup that allows SCU for an unconstrained subject in a virtual reality environment. To this end, we have implemented a portable and fast analog saccade detector inside a head mounted display (HMD) that allows subjects to interact with a 3-D environment, without restrictions on head movement. In addition, this setup also allows more precise positional tracking of gaze direction by means of a separate video based eye-tracker mounted inside the HMD. This is important for our application since we are not only interested in the onsets of saccades but also want to measure absolute eye positions with good accuracy in order to establish how far the changed object was from the fixation points prior and subsequent to the saccade.

Outline of the paper. The remainder of the paper is organized as follows. Section 2 briefly reviews current eye tracking techniques from the perspective of suitability for saccade contingent scene updating in Virtual Reality. In Sec. 3 we give an overview of our system and describe the method of saccade onset detection in detail. Sec. 4 presents experiments that demonstrate the feasibility of the proposed approach. Finally, Sec. 5 concludes our findings and discusses the approach from a broader perspective.

ETRA'02 New Orleans Louisiana USA

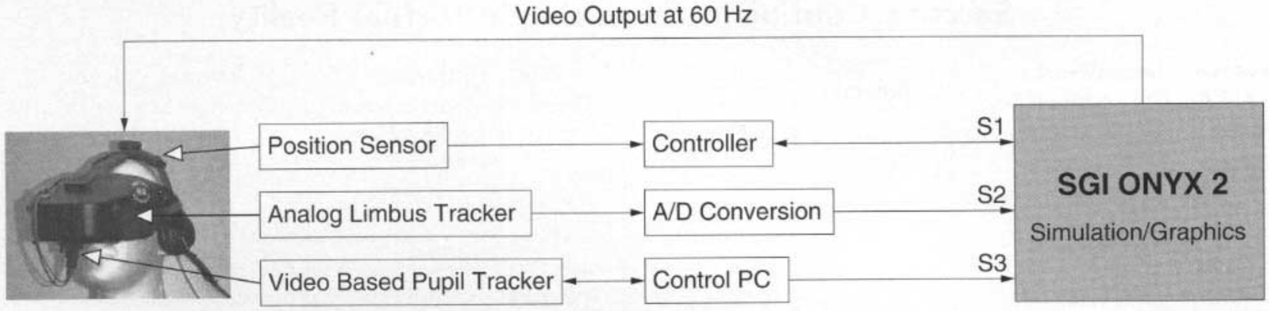

Figure 1: System overview, see text for details.

2 A Brief Review of Eye Tracking Techniques

In order to be able to make saccade contingent scene updates in virtual reality as motivated above, we have to be able to detect the onsets of saccadic eye movements with a very low latency. We will now briefly review some basics about saccade physiology and then discuss several eye tracking techniques from the perspective of their accuracy, their suitability for low latency saccade onset detection, and the feasibility of integration into a head mounted display. A more comprehensive review of eye tracking techniques can be found in [Duchowski and Vertegaal 2000]. We will argue that a combination of two complementary eye trackers is best suited for our needs which will lead to the specific design described in the next section.

Saccade basics. The duration of a saccade typically depends on its amplitude, i.e. the visual angle covered by the saccade. For a more detailed discussion the reader may refer to [Becker 1991]. A good first order approximation of saccade duration D as a function of saccades amplitude A is given by the linear relation:

$$D = D_0 + dA\,, \tag{1}$$

where D_0 is a minimal time typically reported to be between 20 and 30 milliseconds, and d is the duration increase per degree and typically of the order of 2–3 milliseconds per degree. Considering a moderately small saccade amplitude of 20 degrees and using $D_0 = 25$ milliseconds and $d = 2.5$ milliseconds per degree, we get a total duration of 75 milliseconds. The velocity profile of a saccade is typically bell shaped with a peak velocity of roughly about 500 deg/sec. For saccades smaller than about 30 degrees the eye may not reach this peak velocity but stay below it. Since our method of saccade onset detection will be velocity based it is important to know how fast the eye reaches a particular velocity. From the data presented in [Becker 1991] we see that the rise time is about 13 deg/sec/msec for all saccade amplitudes, so the eye should reach a velocity of 100 deg/sec in less than 10 ms.

Electro-oculography. In Electro-oculography electrical signals are recorded from the skin around the eyes with a number of electrodes giving an estimate of the eyes' position inside the head. Although still in use, this method is less accurate than modern video based eye trackers but can give a very high temporal resolution due to the analog output. We undertook preliminary experiments with this method but it seemed that Limbus tracking described below has lower noise and is less intrusive so that we abandoned this method.

Scleral contact lens. In this method the subject wears a contact lens, often with an embedded coil, whose position is measured with respect to a fixed laboratory reference frame. While this method is very fast and very accurate, it is also the most intrusive. In addition, it is not suitable for Virtual Reality applications since the measurement apparatus is too bulky to be mounted inside a HMD.

Limbus tracking. The limbus is the border between the sclera and iris. Across this border there are obvious differences in reflective properties and this can be exploited to estimate the eye's position with optical methods. In a typical setup, an infrared LED illuminates the eye from below and two photo-diodes mounted on spectacle frames measure the amount of light scattered back from the left and right side of the eye. Differences in light received by the photo-diodes allow to estimate the horizontal gaze position. (Typically, a second set of photo-diodes measures light scattered from the upper and lower portion of the second eye to estimate vertical gaze position.) This method is not as precise as modern video based approaches but has the advantage of low latency and high temporal resolution since the analog voltages from the photo-diodes can be sampled and processed without significant temporal delays.

Video-based tracking. Modern video based eye trackers have a video camera pointed at the eye. Often the eye is illuminated with infrared light and the tracking algorithm tries to extract the corneal reflection of the infrared light source and the pupil position from the video images. This method gives fairly good accuracy, better than Limbus tracking although not as good as contact lens methods but suffers from poor temporal resolution. The temporal resolution is limited by the frame rate of the video camera and due to the overhead in capturing and analyzing the video images the eye position measurements are available only with substantial latencies that make saccade contingent updates currently impossible. The use of high speed video cameras (> 100 Hz) can alleviate these problems but these are still too bulky to be built into a head mounted display.

In summary, no single method satisfies our needs for a high accuracy tracker with the capability of low latency saccade detection that can be built into a head mounted display. Other systems, such as the Dual Purkinje Image Tracker, can be successfully used for SCU [Hayhoe et al. 1998; Henderson and Hollingworth 1999; Verfaillie and De Graef 2000], but require a bite bar and are too big to be used inside an HMD and have a very limited range of operation (+/- 7 deg). In order to achieve SCU for a freely moving subject inside a HMD we decided to combine two complementary eye tracking techniques in our system. We chose to integrate a standard video based tracker giving high accuracy with a Limbus tracker for the quick detection of saccade onsets into a single head mounted display.

3 System Organization

3.1 Hardware

The system was developed as part of the University of Rochester National Institute of Health (NIH) Resource for Neural Models of Behavior. A graphical overview of the system is shown in Fig. 1. The backbone of our system is a SGI ONYX 2 workstation with four 250 MHz MIPS processors and two InfiniteReality2 graphics boards from Silicon Graphics. It is responsible for visual simulation of different virtual reality environments including a driving simulator and haptic simulation with force feedback devices, that allows the user to touch virtual objects and physically interact with them. The virtual environments are rendered in a head mounted display (V8 by Virtual Research Systems) with two LCD displays of 640 x 480 pixels (50 deg field of view) at a frame rate of 60 Hz. The head mounted display is equipped with a magnetic head tracker and two complementary eye trackers whose outputs are fed back to the SGI workstation.

The head tracker is a magnetic tracking device (Fastrak by Polhemus Inc.) that senses head position and orientation. These measurements are sent to the workstation over a serial interface.

The first eye tracker is a limbus tracker (ASL model 201 by Applied Science Laboratories) that we have fitted into a pair of swim goggles that can be worn underneath the head mounted display. The original 201 model has two sensors, one for detecting horizontal movements in one eye and one for the detecting vertical movements in the other eye. In our setup we can only use one sensor detecting horizontal movements, because place for the second sensor is taken up by the video based tracker analyzing the other eye. Hence, our setup is not designed to detect the onsets of vertical saccades. Figure 2 shows an image of the sensor and how a subject is wearing it underneath the HMD. The analog voltage output of this sensor is sampled and converted to a digital signal at a rate of 1.25 kHz, and the digital signal is sent to the work station via a second serial connection with 112 kilo bits per second.

The second eye tracker is a standard video based tracker (ASL 501 by Applied Science Laboratories) controlled from a standard desktop PC. It gives relatively good accuracy (~1 deg) at the price of high latency since the video input is sampled at 60 Hz only and there is a delay of about 50 ms before the eye position data is output. This data is sent to the work station at 60 Hz over a third serial connection.

3.2 Saccade Contingent Scene Updating

The time course of events during saccade detection is illustrated in Fig. 3. Subjects may initiate saccades at any time but the graphics engine runs at a fixed rate of 60 Hz so that scene changes can only be made at discrete points in time when a frame is being rendered, i.e. every 16.67 milliseconds. The eye position data from the Limbus tracker, however, arrives at a much faster rate of 1.25 kHz, i.e. individual samples are 0.8 ms apart. Our algorithm for saccade contingent updating works as follows. Just before a new frame is about to be rendered the latest eye position data from the Limbus tracker that has arrived since the last frame is polled and analyzed. Eye velocities are computed from the eye position data and are used to detect the signature of a saccade onset. This way, if a saccade onset is detected the scene can quickly be updated just before it is being rendered again. Our saccade onset detection is based on a simple velocity threshold. If the difference in subsequent eye position readings indicates that the eye is moving at an absolute velocity that is greater than 100 deg/s for 5 ms, this is interpreted as an indication that a saccade has started. Requiring the velocity to exceed 100 deg/s for 5 ms makes the method robust to noise and avoids triggering the detector for very small saccades where the eye barely reaches 100 deg/s. Note that the Limbus sensor only measures the horizontal component of the eye's velocity, which should make it blind to most vertical saccades.

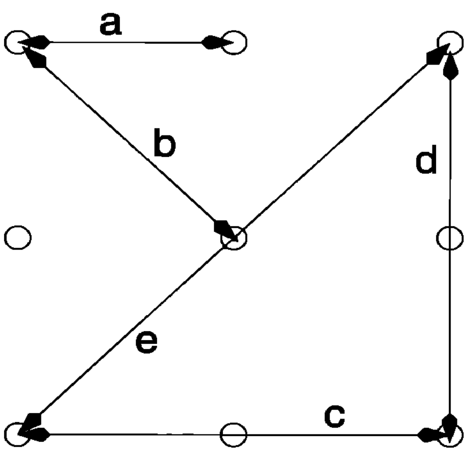

Figure 4: For testing the limbus tracker's performance inside the HMD we had subjects make saccades across a calibration pattern. We tested the shown types of saccades covering visual angles from 6.22 to 17.60 degrees (see text).

The total time between the start of a saccade and when a contingent scene change is visible in the display is the sum of a number of contributions:

$$T = T_{\text{rise}} + T_{\text{transm.}} + T_{\text{detect}} + T_{\text{timing}} + T_{\text{draw}} + T_{\text{display}} . \quad (2)$$

T_{rise} is the time it takes the eye to reach the threshold velocity and should be about 10 ms. $T_{\text{transm.}}$ summarizes delays due to analog to digital conversion and transmission of data over the serial connection. It can be expected to be below 2 ms. T_{detect} is the time required to detect the onset of a saccade in the signal. Since we require the estimated eye velocity to exceed 100 deg/s for 5 ms T_{detect} is 5 ms. T_{timing} is a variable delay that is determined by the relative timing of saccade detection with respect to the fixed times of scene rendering. It can be expected to be between 0 and 12 ms. T_{drawing} is the time it takes the graphics engine to cull and draw the scene and then to display it. This number depends a little on scene complexity. For a typical application T_{drawing} is around 24 ms. Finally, T_{display} is the time it takes the LCD displays in the helmet to switch state. We measured a time of 10 ms (see below). Hence, we can estimate the total end-to-end latency to be between 51 ms and 63 ms. This indicates that the system should be able to change the scene before the eye comes to rest for horizontal saccades greater than roughly 15 degrees.

4 Experiments

4.1 Feasibility Tests

Our initial experimentation was concerned with two main questions: First, how well does the saccade detector work underneath the HMD? Second, is the system latency in this setup small enough to detect a saccade and make an on-screen change while the eye is still in motion? To test the limbus tracker's effectiveness inside the HMD, the following general procedure was used: Subjects were calibrated for both the limbus and video based tracker and then presented with a sequenced test pattern. Four female subjects, all Uni-

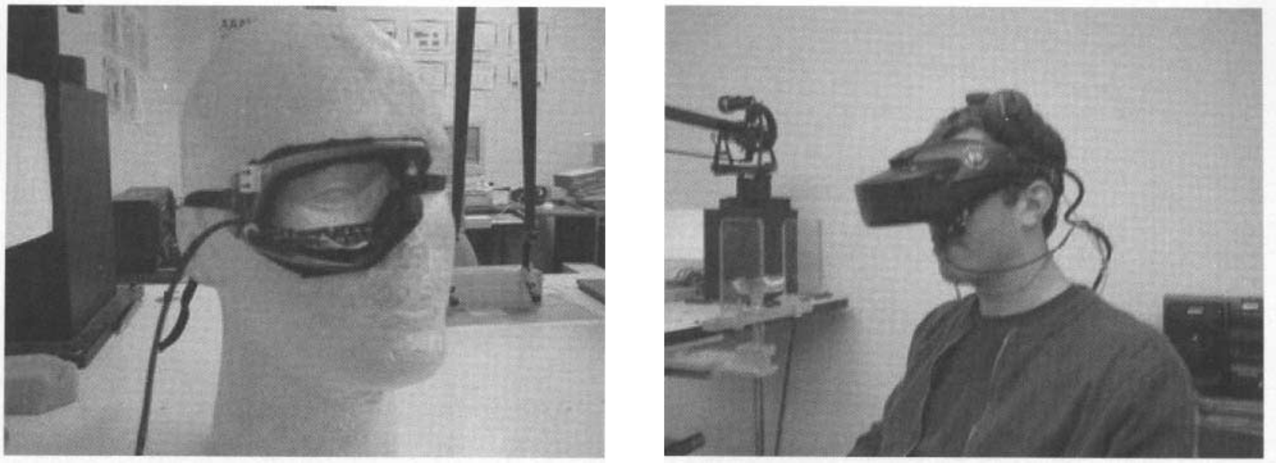

Figure 2: **Left:** Limbus eye tracker mounted in swim goggle frame. **Right:** Subject wearing Limbus tracker and HMD with built in video based tracker.

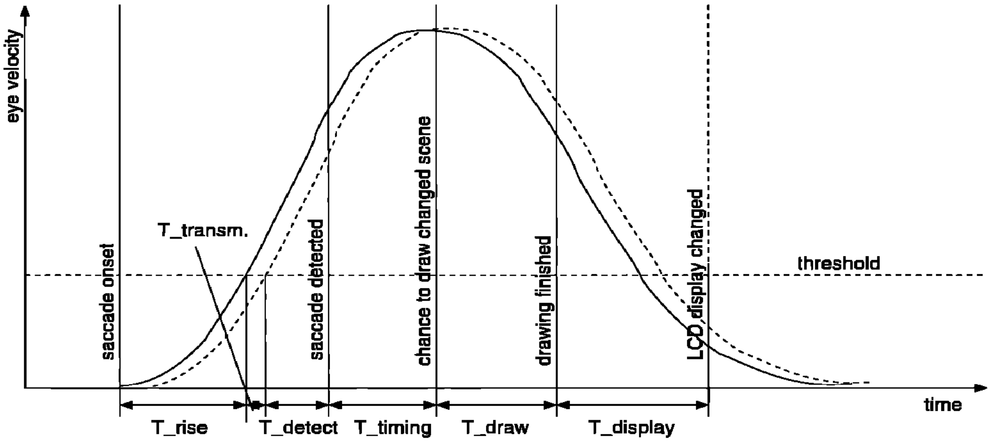

Figure 3: Course of events during saccade detection. The solid bell shaped curve depicts the absolute horizontal eye velocity as a function of time, the dashed curve shows when this data can be accessed in the computer. Individual contributions to total system latency are not drawn to scale.

versity of Rochester students, participated in this study and received monetary reimbursement. Two additional male subjects, also University of Rochester students were used but the data was thrown out due to errors in the video based tracking. By presenting a predetermined sequence of fixation points we had explicit knowledge of what the subjects eyes should be doing during the entire experiment. The sequenced pattern worked in the following manner: A test pattern consisting of a 3x3 grid of the numbers one through nine evenly spaced at 100 pixels or 6.22 degrees of visual angle was presented to the subjects in a piecemeal fashion, such that only one number appeared at a time. Subjects were instructed to fixate on the number presented and then immediately shift their gaze to the next number that appeared in the sequence. Subjects made the following 14 discrete eye movements for this sequence (see Fig. 4): six 6.22 degree horizontal saccades (a), two 8.79 degree diagonal saccades (b), three 12.44 degree horizontal saccades (c), one 12.44 degree vertical saccade (d), and two 17.598 degree diagonal saccades (e). The decision to also test vertical saccades (d) despite the system s design for measuring only horizontal eye velocity was motivated by earlier experience with the standard version of this tracker that suggested substantial cross-talk between the horizontal and vertical direction. Once the subjects completed looking through the sequenced pattern the data was analyzed for the number of saccades detected as well as the latency between the onset of a saccade and the point at which it was detected by our algorithm. Each subject looked through the pattern a total of three times. All subjects were sitting down during the experiment. While the subjects were not explicitly told to limit their body movements, they remained still for the most part.

Trial	CS	BD	MD	AS
1	100	92.9	100	92.9
2	92.9	92.9	100	71.4
3	100	92.9	100	64.3
total	97.8	92.9	100	76.2

Table 1: Saccade detection rate in percent.

Amplitude (deg)	#	Duration (ms)	horiz. velocity (deg/s)
6.22 horiz.	30	44.2	308.0
8.80 diag.	10	48.6	341.9
12.44 horiz.	14	57.3	421.1
12.44 vert.	4	55.4	221.3
17.60 diag.	10	69.1	406.0

Table 2: Measurement of some saccade characteristics. The column marked # shows the number of trials performed. Note that the sensor measures a substantial velocity for the vertical eye motion although being operated in the horizontal direction. This cross-talk effect is common in limbus trackers.

Table 1 summarizes the results. It shows the percentage of detected saccades for four different subjects. The tracker performs very well. With the exception of subject AS the limbus tracker detected all the horizontal and diagonal saccades. The only saccades being missed are vertical saccades which the system is not designed to detect. It is interesting to note, however, that some of the vertical saccades are detected nevertheless, which can be attributed to the simple design of limbus trackers that often leads to cross-talk between horizontal and vertical eye movements. The somewhat worse performance for subject AS may be due to a bad calibration. Averaged over all subjects, the tracker detects 92% of the total saccades. We found that sometimes individuals blinked or made additional eye movements that were detected as saccades. In these instances the data was ignored and it was generally assumed that the eye movement that occurred directly after the new visual stimulus was presented was the saccade of interest. We will discuss the detection of false positives due to, e.g., blinks in more detail below.

Beyond detecting the onsets of saccades with low latencies we can also use the setup to measure durations and velocities of saccades (with the limitation of only having access to the horizontal component of the movement). Table 2 shows the average durations and peaks of the horizontal velocity components for a number of saccades from 4 different subjects. Note that the vertical eye movement nevertheless creates a substantial velocity signal in the limbus tracker which is aimed at only detecting horizontal movements.

Regarding relative latency of saccade onset detection for the two eye trackers, Fig. 5 compares the output of the limbus tracker and the video based tracker. Individual data points for the limbus tracker are 0.8 ms apart (light gray curve), while the video based tracker only gives one new data point each 16.7 ms (medium gray curve) and has substantial latency. Our simple velocity threshold algorithm estimates the absolute eye velocity (dark curve) from the limbus data and detects the saccade about 18 ms after the first signal deflection from its resting state. We also tried to detect the saccade onset on the basis of the video based tracker's signal. To this end we used the same velocity threshold criterion. In the example shown in Fig. 5 detection would occur at time $t = 400$ which is about 26 ms later than for the limbus tracker. We systematically tested the difference in detection latency for the limbus tracker and the video based tracker. We analyzed 15 saccades and found that the limbus tracker detects saccade onsets between 27 and 50 ms earlier with an average advantage of 37 ms.

4.2 Total System Latency

While our initial tests of system latency gave a good evaluation of relative latency of the limbus tracker and the video based tracker it does not allow to measure absolute latencies. Our estimate from Sec. 3 suggested an end-to-end latency between 51 and 63 ms. We verified our analysis by measuring the end-to-end system latency using the following setup: A Tektronix Dual Channel Oscilloscope model TDS 3032 was set up such that it could monitor the input and output of the system. One channel directly monitored the raw analog voltage output of the limbus tracker while the second channel measured the voltage of a photo-sensor, a PIN photo-diode running in current mode with a convex lens and fully shielded, placed on the left LCD eye piece. An example measurement is shown in Fig. 6. To measure the system latency, an artificial saccade was triggered by the experimenter by waving a finger once over the limbus tracker at a velocity that was sufficient to trigger the saccade detector. When a saccade was detected the computer would change the color of the LCD screen from white to black. By monitoring the two voltages the point at which the limbus tracker's output voltage began to deflect can be compared with the time when the LCD actually changes state from white to black. Since the finger movement produced a different and varying rise of the velocity signal very different from a real saccade we focused on the time between onset detection by the algorithm and change of the LCD display. (In future experiments we will eliminate the rise time issue by switching a light in front of the limbus tracker.) We expected this time, the sum of T_{timing}, T_{draw}, and T_{display}, to be between 34 and 46 ms. We performed 10 measurements and found latencies in the range of 33.5 to 45.8 ms, which is in excellent agreement with our prediction. To estimate the end-to-end latency from the true beginning of a saccade to the end of the display change we have to add about 17 ms ($T_{\text{rise}} + T_{\text{transm.}} + T_{\text{detect}}$), which gives a worst case end-to-end delay of about 63 ms. Thus, assuming the correctness of our estimates for T_{rise} $T_{\text{transm.}}$, and T_{detect}, our setup allows to change the scene before the eye comes to rest if the saccade amplitude is greater than about 15 deg. Note that if using a CRT as a display the update would be substantially faster.

4.3 Saccade Detection with Head Motion

With the initial tests completed a new segment of testing was begun concerning the issue of saccade detection while the head was in motion. The same sequenced calibration pattern procedure was used as above however subjects were given the additional instruction that they should move their heads back and forth from left to right while the sequenced pattern was presented. Subjects were told to go at a rate of approximately one head turn per second. Three trials were performed with two new subjects. JA participated in two trials and EG in one.We measured horizontal head rotations of ± 14 deg amplitude with absolute angular velocities between 10–20 deg/s. A second test identical to the first, except subjects were instructed to look up and down, was also run. We measured vertical head rotations of ± 13 deg with absolute angular velocities covering the range from 8–45 deg/s.

Results for the horizontal portion were rather positive as detection neared 86% which is close to the performance without head movement.Results for the vertical motion section is only 69%. A reason for the inferior result for the vertical head movements may be that the HMD tends to significantly shift its position during vertical head movements of the range we considered here, possibly de-calibrating the limbus tracker worn underneath.

4.4 False Positives

Three final tests were conducted to look at the rate of false positives in saccade detection. Three distinct eye movements, vestibulo-ocular reflex (VOR, a reflex that rotates the eyes to compensate for head movements), smooth pursuit, and eye blinks could possibly fool the saccade detection algorithm into marking a saccade that is not actually occurring. We addressed this question in three experiments.

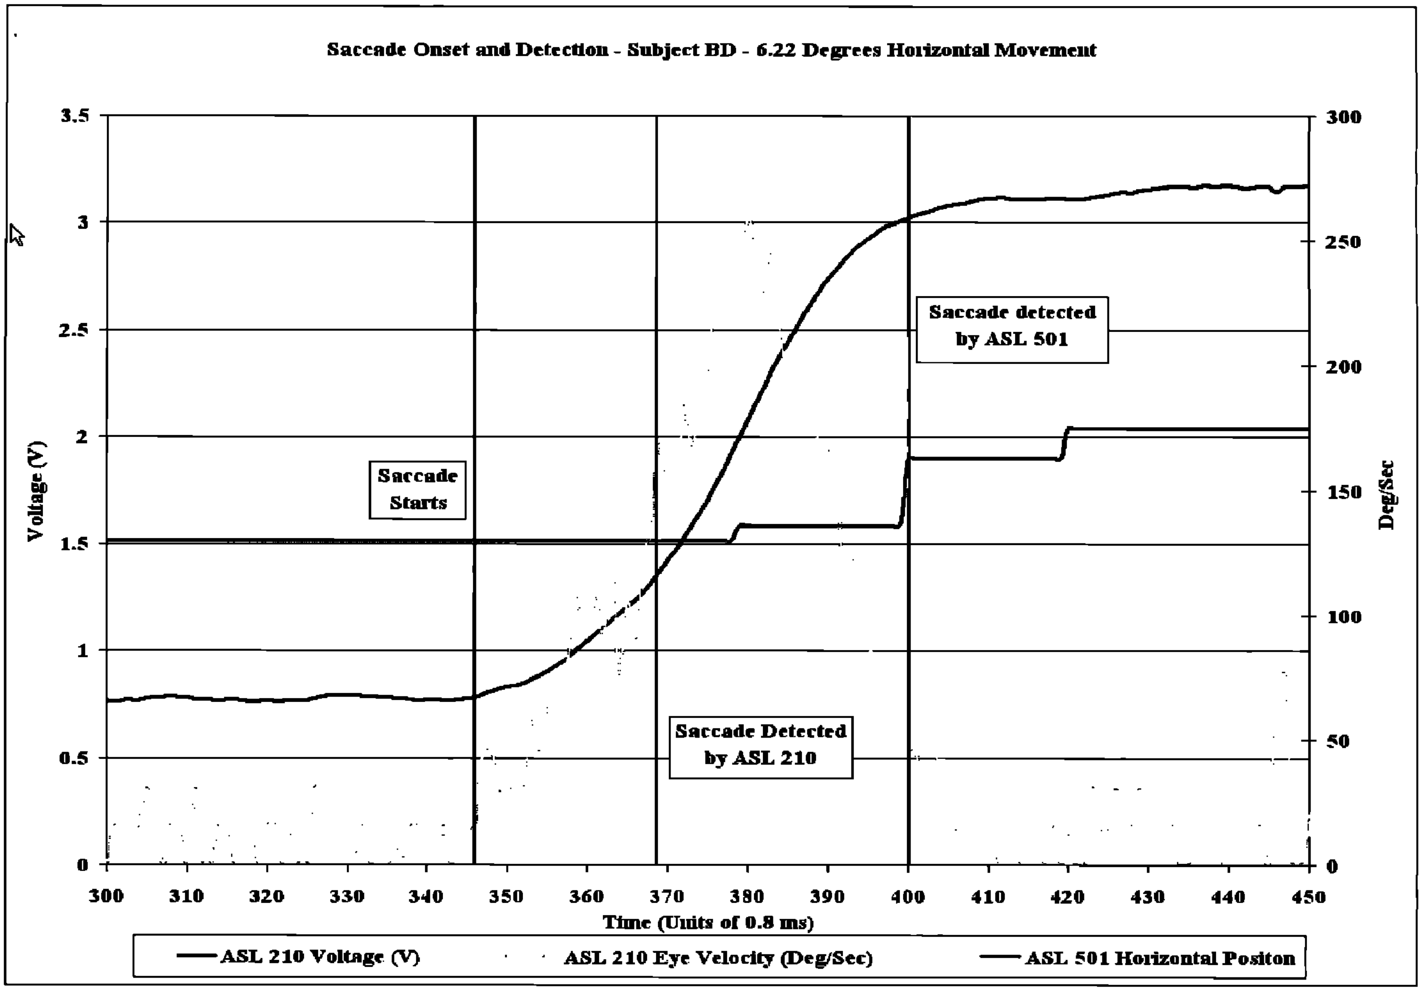

Figure 5: Comparison between signal latency of the limbus tracker and the video based tracker. The onset of the saccade can be seen in the limbus tracker's signal (raw voltage: dark curve, eye velocity: light gray curve) much earlier than in the signal from the video based tracker (medium gray curve) which is sampled at only 60 Hz. In this example the eye velocity first exceeds 100 deg/s roughly 10 ms after the first signal deflection in accordance with our estimate based on saccade physiology from Sec. 2. The saccade is detected a couple of milliseconds later.

VOR. In the VOR portion subjects were calibrated with both eye trackers and instructed to move their heads back and forth from left to right and to always keep their eyes fixated on a stationary object, (the back of a virtual car approximately 3 deg wide and 2 deg tall). Subjects rotated their head left and right with amplitudes between ± 15-20 deg and average angular velocities of 11 deg/s. The data was examined for any saccades detected. Although subjects were instructed to keep the object fixated an analysis of their eye movements from the video based tracker and video tapes of the subjects' eyes during performance revealed a number of real saccades that had triggered the saccade detector. For an estimate of only the *false* positives we had to sort out these correctly detected saccades. We tested three subjects for 27s , 27s, and 48s, respectively. We found that all saccade detections announced by the detector corresponded to real saccades or eye blinks and we did not find a single false alarm based on VOR motion in this data, suggesting that the method is quite robust with respect to false alarms of this kind.

velocity (deg/s)	3.2	15.8	33.6
false alarms per 10s	0.8	0.2	0.5

Table 3: Average false alarm rate during smooth pursuit.

Smooth pursuit. To test for false positives during smooth pursuit subjects were again calibrated with both trackers and then instructed to track an object presented in the display. In these trials the same car as above was used but now it moved from side to side and faced the subject lengthwise. The car's extent at this aspect was approximately 4 degrees horizontally and 2 degrees vertically. The car moved from side to side at 3 different velocities that were tested separately (3.2 deg/s , 15.8 deg/s , and 33.6 deg/s). Single trials lasted on average 35, 20, and 15 seconds for the three object velocities. A similar methodology to the VOR portion was used for the data analysis. The results are summarized in Tab.3. It shows the average number of false alarms observed in 10 seconds of recording

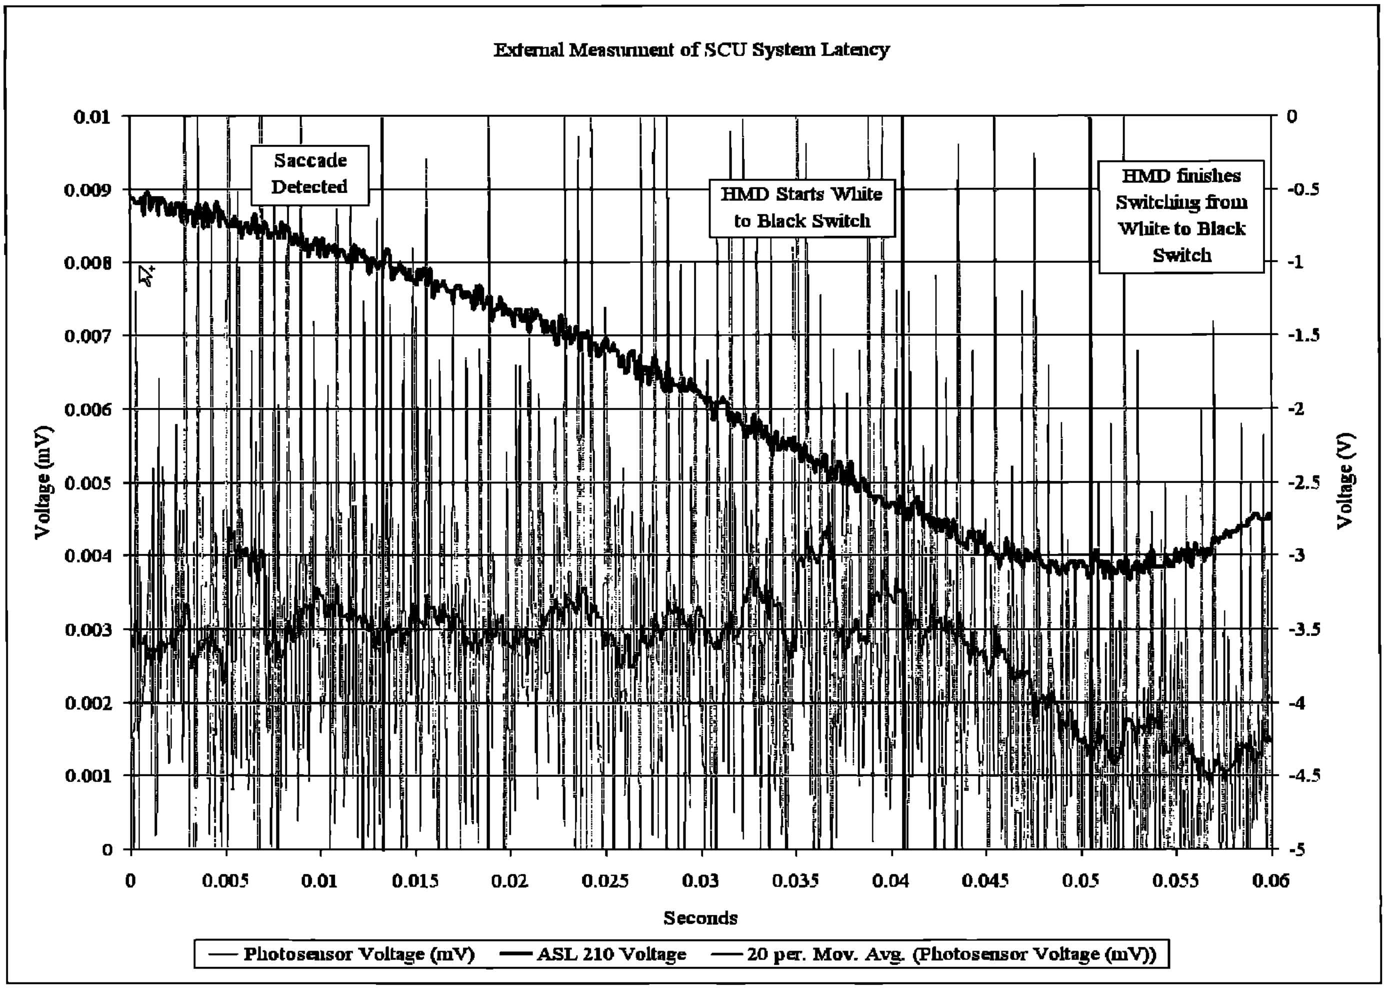

Figure 6: Example result of end-to-end system latency test. The dark curve is the raw signal recorded by the limbus tracker. The saccade detection was triggered by waving a finger in front of the sensor which gives a somewhat slower rise time of the signal as compared to a saccade. The leftmost vertical bar marks the time when the saccade onset was detected. The light gray data shows the voltage of the photo-sensor that measured when the HMD display changed color. The medium gray curve is a 20 data point wide running average of this data. The LCD display reacts somewhat sluggishly, taking about 10 ms to change the color. Time between detection of saccade and end of display change is 38 ms which falls well in the predicted range of 34–46 ms for $T_{\text{timing}} + T_{\text{draw}} + T_{\text{display}}$.

for the three velocities. The data is averaged from 4 subjects participating in the study. As can be seen, the false alarm rate during smooth pursuit movements is generally below one in 10 seconds, which seems reasonably good. Surprisingly, however, the rate of false alarms does not depend on the velocity of the tracked object, contrary to our expectations. This suggests that at least some of the false alarms might have a cause other than the smooth pursuit movements.

Blinks. Finally, the effect of blinks on the limbus tracker was examined. Subjects were calibrated for both eye-trackers and then asked to blink 20 times in a row separated by roughly one second intervals while staring straight ahead. In this portion of the study the subjects EG and JA participated. The data files were recorded and examined for the number of saccades detected compared against the ASL 501 data. For both subjects all 20 blinks triggered the detection of a saccade. For both subjects the signature of blinks in the limbus tracker signal was stereotypical in that it always deflected the voltage of the ASL 210 tracker in the same direction (negative for JA and positive for EG). According to the manufacturer this is a typical signature of blinks.Distinguishing blinks from saccades at their onset is very problematic. It should be pointed out, however, that for the kind of change blindness applications we are interested in, making a scene change during a blink is also an accepted technique for studying change blindness [O Regan et al. 2000]. Hence, our current inability to distinguish saccade onsets from blink onsets does not pose a problem for the applications we are most interested in.

5 Conclusions

We have presented a method for saccade contingent updating (SCU) in virtual reality. Our system integrates two complementary eye trackers into a single head mounted display (HMD). The first is a standard video based tracker giving high accuracy at the price of a high latency. The second is an analog limbus tracker that gives low latency signals at the cost of less accuracy. While researchers have used SCU earlier, we have shown that it can be done in virtual reality, i.e. while wearing a HMD, and without the need to constrain subjects' heads. However, we found that vertical head movements seemed to easily de-calibrate the limbus tracker and this point needs to be studied more carefully.

We performed a number of experiments to test the system's performance in detecting saccades with low latencies. Our tests have proven the principal feasibility of saccade contingent updating in a virtual environment. While the method is not perfect it has very few missed detections and low rates of false positives during vestibulo-ocular reflex and smooth pursuit. Eye blinks, however, tend to be systematically mistaken for saccades. For our intended applications of the method, however, this may not turn out to be a severe limitation, since eye blinks are also very effective in masking visual transients [O'Regan et al. 2000].

While it is safe to say that we can make saccade contingent updates for horizontal saccades greater than about 15 degrees we cannot currently guarantee this for smaller saccades. This means, that for smaller saccades we may trigger a change that only becomes visible when the eye has already finished the saccade. However, we think that this is not a critical problem for our applications for two reasons. First, saccadic suppression may be active until after the end of a saccade [Stevenson et al. 1986]. Second, we can easily structure the environment to encourage big saccades at certain times in the task by controlling the distances between relevant objects in the scene.

In conclusion, we were able to demonstrate that saccade detection and contingent updating is not limited to the realm of bite bars and chin rests but can be accomplished reasonably well for a freely moving subject immersed in virtual reality. Our setup could have potentially liberating effects for experimenters (and their subjects) interested in making saccade contingent updates in settings that do not limit the subjects' head movement and thus allow a more immersive environment.

While our main motivation now is to exploit the new capability of saccade contingent updating for our studies of human perception and attention [Shinoda et al. 2001], we would like to briefly discuss an interesting avenue for future research on the eye tracker itself. Since the velocity profiles of saccades often tend to be symmetric for saccades that are not too big, it would be tempting to try to predict the target location of a saccade while the eye is still "in flight". If this can be done soon enough one can still make a change exactly where the eye will "land". If feasible, this approach might find interesting applications such as seamless gaze-contingent level-of-detail computer graphics for virtual reality.

Acknowledgments

We wish to thank Ted Twietmeyer for assistance in evaluating the end-to-end system latency. This work was supported by NIH Resource Grant P41 RR 09283.

References

BECKER, W. 1991. Saccades. In *Vision & Visual Dysfunction Vol 8: Eye Movements*, R. H. S. Carpenter, Ed. CRC Press, 95–137.

DUCHOWSKI, A., AND VERTEGAAL, R., 2000. Eye-based interaction in graphical systems: Theory and practice. Course notes, available at http://www.vr.clemson.edu/eyetracking/sigcourse.

HAYHOE, M., BENSINGER, D., AND BALLARD, D. 1998. Task constraints in visualworking memory. *Vision Research 38*, 125–137.

HENDERSON, J., AND HOLLINGWORTH, A. 1999. The role of fixation position in detecting scene changes across saccades. *Psychological Science 10*, 438–443.

O'REGAN, J. K., DEUBEL, H., CLARK, J. J., AND RENSINK, R. A. 2000. Picture changes during blinks: Looking without seeing and seeing without looking. *Visual Cognition 7*, 191–211.

SHINODA, H., HAYHOE, M., AND SHRIVASTAVA, A., 2001. Attention in natural environments. Vision Research, special issue on Eye Movements and Vision in the Natural World. (in press).

SIMONS, D. J. 2000. Current approaches to change blindness. *Visual Cognition 7*, 1–15.

STEVENSON, S. B., VOLKMAN, F., KELLY, J., AND RIGGS, L. 1986. Dependence of visual suppression on the amplitudes of saccades and blinks. *Vision Research 26*, 1815–1824.

VERFAILLIE, K., AND DE GRAEF, P. 2000. Transsaccadic memory for position and orientation of saccade source and target. *Journal of Experimental Psychology: Human Perception and Performance 26*, 1243–1259.