
Heuristic reduction of gyro drift in vehicle tracking applications

Johann Borenstein and Lauro Ojeda

University of Michigan
2260 Hayward Street
Ann Arbor, MI 48108, USA
E-mail johannb@umich.edu, lojeda@umich.edu

Abstract: This paper pertains to the reduction of measurement errors due to drift in rate gyros used for tracking the position of moving vehicles. In these applications, gyros and odometry are often used to augment GPS when GPS reception is unavailable. Drift in gyros causes the unbounded growth of errors in the estimation of heading, rendering low-cost gyros almost entirely useless in applications that require good accuracy for more than just a few seconds or minutes. Our proposed method, called “Heuristic Drift Reduction” (HDR), applies a unique closed-loop control system approach to estimate drift in real-time and remove the estimated drift instantaneously from the gyro reading. The paper presents results of experiments, in which a gyro-equipped car was driven hundreds of miles on highways, rural roads, and city streets. HDR reduced the average heading error over all of these drives by one order of magnitude.

Keywords: gyro, drift, vehicle tracking, heuristic

Biographical Notes:

Johann Borenstein received the B.Sc., M.Sc., and D.Sc. in Mechanical Engineering in 1981, 1983, and 1987, respectively, from the Technion – Israel Institute of Technology. He currently holds the rank of Research Professor and has been the Head of the Mobile Robotics Lab at the University of Michigan (USA) for 21 years. He has nine patents, a book, and over 150 other publications on mobile robotics. His areas of expertise include mobile robot (and human) navigation, obstacle avoidance, and innovative mobile robot mechanisms. He is or was the PI on numerous large projects funded by DOE, DARPA, NASA and private companies.

Lauro Ojeda has a Bachelor Degree in Electronic Engineering and a Master of Science Degree in Electronics from the Army Polytechnic School - Ecuador. He is currently a research investigator at the Mobile Robotics Laboratory of the University of Michigan. His research interest is in the area of sensor fusion and inertial navigation systems for mobile robots and humans. He has worked on several projects sponsored among others by DARPA, DOE, and NASA.

1 Introduction

For almost all outdoor land navigation tasks, GPS is the obvious sensor of choice. However, it is well known that dense tree coverage or so-called “urban canyons” can occlude satellite signals. In military applications intentional jamming of GPS radio signals is also a concern. Indoors or underground, GPS is unavailable altogether. Indoor vehicle tracking applications and outdoor applications, in which GPS outages are anticipated, require therefore non-GPS solutions, at least some of the time.

For the vehicle tracking system of this paper, we assume that the vehicle is moving on near horizontal terrain for most of the time and therefore we used a single one-axis gyro to measure the yaw-rate, from which relative changes in heading can be computed.

The output of a gyro is rate-of-turn, ω . In vehicle tracking applications, one is usually interested in heading, which can be computed from ω by integrating the output signal numerically. The numeric integration has a tendency to cause errors due to “bias instability,” more commonly referred to as “drift.” Drift is produced when small, slow-changing deviations from the correct signal are integrated with respect to time. The highly undesirable result of drift is that the error of the computed output—relative heading—increases continuously and without bound.

Gyros are also sensitive to changes in temperature, which incur slow-changing deviations, just like drift does. Our proposed drift reduction method counteracts all slow-changing errors regardless of whether they were caused by the physical phenomena of drift or temperature sensitivity. For that reason, throughout this paper, we lump all slow-changing error sources together and call them collectively “drift.”

Yet another source of errors in computing heading is high-frequency noise in the gyro’s output signal. The accumulated error due to the integration of this signal is called “Angle Random Walk” (ARW). However, in the computation of heading from the gyro’s rate of turn measurements, errors due to ARW are relatively small, since the ARW’s average is about zero.

Before explaining our proposed drift reduction method in detail in Section 2, we discuss here briefly several others approaches. The most common method for reducing the effects of gyro drift is to integrate data from an Inertial Measurement Unit (IMU) and GPS [Cavallo, et. al. 2005] [Mohinder, et. al., 2002], [Grejner-Brzezinska, et. al. 2006]. Another method integrates data from a magnetic compass as well [Cho, et. al. 2003]. The main drawback of these approaches is that they require GPS and/or magnetometer data, which are not available all the time. Basnayake et al. [2005] proposed a method that makes use of available maps and uses map matching techniques for further enhancement. Some other methods have been proposed that try to improve accuracy by find a mathematical model for bias errors [Paniit and Wwibang, 1986][Chen, X, 2004]. However, these techniques have limited applicability and can only estimate the deterministic part, if any, of the bias drift.

Our proposed method, called “Heuristic Drift Reduction” (HDR), exploits the fact that many surfaces, on which vehicles travel are man-made and oftentimes have straight line features. At any moment, the HDR method estimates the likelihood that the vehicle is driving along a straight line. If that likelihood is high, HDR applies a correction to the gyro output that would result in a reduction of drift if indeed the vehicle was driving along a straight line. If HDR assesses that the vehicle is not traveling straight, then it does nothing. This way, and on an abstract level, HDR uses landmarks (i.e., man-made straight-line features), but there is no requirement that the location of these landmarks be known in advance.

The remainder of this paper is organized as follows. Section 2 explains the basic HDR method in the context of ideal conditions. Section 4 discusses how real conditions differ from ideal conditions, and provides several enhancements to the HDR algorithm to help

Heuristic reduction of gyro drift in vehicle tracking applications

cope with real conditions. Section 5 shows experimental results obtained with HDR applied to a gyro in a vehicle tracking system. Conclusions are presented in Section 5.

2 Heuristic Drift Reduction

Suppose a vehicle is driving straight forward. Moving straight forward, the output of the gyro should be exactly zero throughout the trip. However, due to drift the actual output is off by some small value, ε . Suppose further that we divided the total travel distance into one-second intervals. Due to the drift error ε , in each interval the rate of rotation computed based on the gyro measurements is thus

$$\omega_{raw,i} = \omega_{true,i} + \varepsilon_0 + \varepsilon_{d,i} \quad (1)$$

where

- $\omega_{raw,i}$ – Rate of rotation measurement. This is the direct output of the gyro.
- $\omega_{true,i}$ – True rate of rotation. In reality ω_{true} is not known or measured, but in our idealized straight-line example $\omega_{true} = 0$.
- ε_0 – Static bias drift, measured immediately prior to a drive.
- $\varepsilon_{d,i}$ – Drift (i.e., bias instability).

Immediately prior to each drive and with the gyro held completely motionless, the static bias drift ε_0 is measured by averaging $T_{bias} = 30$ seconds worth of gyro output. The value for T_{bias} depends on the quality of the gyro and can be estimated by using the Allan Variance analysis [IEEE, 1996]. T_{bias} is also called “bias time.”

During the drive, ε_0 is subtracted from every reading of $\omega_{raw,i}$:

$$\omega'_{raw,i} = \omega_{raw,i} - \varepsilon_0 = \omega_{true,i} + \varepsilon_{d,i} \quad (2)$$

Then, the new heading ψ_i is computed

$$\psi_i = \psi_{i-1} + \omega'_{raw,i} T_i \quad (3)$$

where

- ψ_i – New computed heading [$^\circ$].
- T_i – Duration of time interval i [sec]. Throughout this paper, $T_i = 1$ sec.

3 The Basic Heuristic Assumption

During straight-line travel, if $\varepsilon_{d,i}$ is positive, then the perceived change of heading in this interval is counter-clockwise (we will call this a “left turn,” for simplicity) and if $\varepsilon_{d,i}$ is negative, then the perceived change of heading is clockwise or a “right turn.”

It is unpredictable whether ε_d will be positive or negative and ε_d may change signs during a drive. However, if ε_d changed signs very often during a drive, then drift errors will partially cancel each other out and the overall error is less severe. Our greater concern is thus for cases where ε_d keeps the same sign for prolonged periods of time and thereby accrues heading errors in the same direction. For the sake of simplicity, we assume in the remainder of this section that ε_d keeps the same sign throughout the drive. In practice, however, this is not a necessary requirement for the HDR algorithm and ε_d may change signs.

If ε_d keeps the same sign throughout the drive, then in each interval T_i the heading error will have the same direction regardless of the unpredictable and ever-changing value of ε_d .

For a straight-line drive of 1,000 seconds and assuming ε_d is positive, there are 1,000 intervals, in which the vehicle tracking system erroneously perceived that it had turned

left (due to the positive drift assumed in this example), and zero intervals, in which it erroneously perceived it had turned right. Noise in the gyro output can blur this sharp ratio somewhat, but as long as the average of the noise is zero there will still be a significantly larger count of perceived left turns than right turns. Moreover, once the absolute drift value exceeds the magnitude of the noise, noise will no longer have any influence on the left turn/right turn count.

More importantly, even a sharp turn will affect this counting scheme only briefly, namely for the duration of the turn. Once the turn is over, there will again be many more perceived left turns, due to the positive drift.

Next we will introduce a practical method that can effectively estimate the magnitude of drift based on the proceeding observations.

3.1 The Basic HDR Algorithm

In order to explain the basic HDR algorithm, we continue to make the simplifying assumption that when a vehicle drives, it moves along straight lines—at least some of the time. In a later section we will introduce enhancements to the basic HDR method that allow us to drop the simplifying assumption of this section and deal effectively with all realistic conditions.

The basic HDR algorithm functions essentially like a closed-loop control system. This is different from most other measuring systems, where signals pass from the sensor to the instrument’s output in open-loop fashion. Figure 1 shows a block diagram of the closed-loop control system implemented in the HDR algorithm.

Our above-stated simplifying assumption means that $\omega_{true}=0$ is correct at least some of the time. When $\omega_{true}=0$, then the only output from the gyro (after subtracting the static bias drift ϵ_0) is ϵ_d . For our closed-loop control system, ϵ_d is a disturbance. Because of the I-controller, and provided the control parameters are properly chosen, the control signal I will track (but with an opposite sign) slow changes of ϵ_d with no offset. That is, $I \cong -\epsilon_d$ under ideal conditions.

The ideal condition $\omega_{true}=0$ is, of course, rarely met. Indeed, ω_{true} can briefly be orders of magnitude larger than ϵ_d , for example, when the vehicle takes a turn. In that case a conventional I-controller does not work well, since it would respond strongly to large values of $|\omega_{true}|$, thereby overwhelming the integrator in the I-controller. To avoid this pitfall, the I-controller should be insensitive to the magnitude of the error signal E . This can be achieved by treating the error signal E as a binary signal that can have only one of two

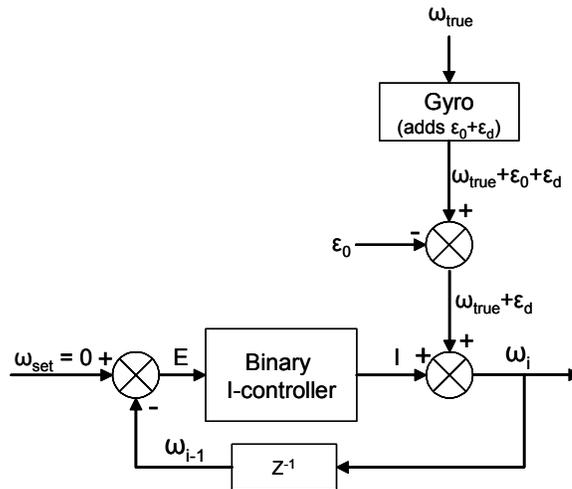


Figure 1: The basic HDR algorithm viewed as a closed loop control system. The binary I-controller is explained in the text.

Heuristic reduction of gyro drift in vehicle tracking applications

values: positive or negative. This way, the integrator I reflects the difference between the numbers of perceived left and right turns, as explained in the proceeding section.

For the implementation of the algorithm we should note that since the setpoint ω_{set} is permanently set to zero, the following is true:

When $\omega_{i-1} > 0$ (a perceived left turn), E is negative

When $\omega_{i-1} < 0$ (a perceived right turn), E is positive

We can now formulate the binary I-controller

$$I_i = \begin{cases} I_{i-1} - i_c & \text{for } \omega_{i-1} > 0 \text{ (a perceived left turn)} \\ I_{i-1} + i_c & \text{for } \omega_{i-1} < 0 \text{ (a perceived right turn)} \end{cases} \quad (4a)$$

where

i_c – Fixed increment [$^\circ/\text{sec}$]

An alternative way of writing Eq. (4) is

$$I_i = I_{i-1} - \text{SIGN}(\omega_{i-1})i_c \quad (4b)$$

where $\text{SIGN}()$ is a programming function that determines the sign of a number. SIGN returns ‘1’ if the number is positive, ‘0’ if the number equals zero, and ‘-1’ if the number is negative.

The next element in the control loop adds the controller output to the raw measurement

$$\omega_i = \omega_{true} + \varepsilon_d + I_i \quad (5)$$

where

ω_i – Corrected rate of rotation [$^\circ/\text{sec}$].

If $I \cong -\varepsilon_d$, as we assume for now to be the case under ideal conditions, in steady state, and because of the closed loop control system, then by substituting $I \cong -\varepsilon_d$ in Eq. (5)

$$\omega_i \cong \omega_{true,i} \quad (6)$$

This result is desirable, since the unknown drift is removed.

4 Realistic Conditions

In order to explain the basic HDR algorithm we assumed that much of the driving was along perfectly straight lines. During ideal straight-line motion, any perceived right or left turn is indeed the result of drift. In reality, however, vehicles don’t drive exactly straight. In Section 4.1 we discuss typical detractions from ideal straight-line driving, such as swaying, curving, and turning. Then, in Section 4.2, we present enhancements to the basic HDR algorithm, aimed at counteracting these detractions.

4.1 Detractions from Ideal Straight-line Motion

4.1.1 Swaying

We call motion that is *intended* to be straight but is not entirely straight “swaying.” Swaying is the result of minute steering corrections that a human driver performs intuitively.

tively in order to stay near the center of a traffic lane. An example for swaying is shown in Figure 2. Errors due to swaying can be reduced significantly by low-pass filtering, as will be explained in Section 4.2.1.

4.1.2 Curving

Curving is the motion along an extended arc. This motion is potentially the least favorable one for the HDR method. That is because extended driving along an arc with a very large radius will be perceived continuously as a turn in the same direction and counteracted just like drift. In the worst case, if the arc curves, say, to the left, while the actual drift is to the right, then the HDR algorithm will try to counteract the perceived left turns and increase I , thereby actually increasing the error caused by drift alone.

This is a particular concern with minor roads that follow natural terrain and consist mainly of curving segments. In those cases, the basic HDR will not work well. Being aware of this limitation, we developed a method for detecting curving motion. This method, described in detail in Section 4.2.3, gradually reduce the effect of HDR as long as the curving motion persists, so as to avoid the introduction of greater errors. The effectiveness of this measure will become apparent in the “Suburbs” experiment listed in Section 5. In that experiment we intentionally drove along a route that comprised mostly of the curving streets found in many typical suburban subdivisions in the United States. Even then HDR reduced heading errors by a factor of 3.4 (see Section 5). This improvement was possible because even curving roads have short straight-line sections, and HDR tracks drift whenever the vehicle is on a straight line section.

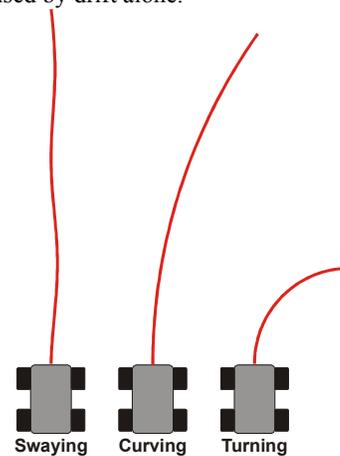


Figure 2: Three types of non-straight motion: Swaying, Curving, and Turning.

4.1.3 Turning

Turning is a sharp but short change of direction. Examples are street corners or intersections of rural roads. Turns are easily identified because the gyro measures much larger rates of rotation during turning than what can be expected as a result from drift. Thus, a simple test can be performed in software: if $|\omega|$ is larger than some threshold, then Eq. 4 can be skipped altogether.

4.2 Refinements to the Basic HDR Algorithm

In this section we discuss several enhancements to the basic HDR algorithm. Together, these enhancements overcome the two main challenges for the HDR method under real conditions: swaying and curving.

4.2.1 Low-pass Filter

Passing raw gyro readings through a low-pass filter is normally not necessary for estimating heading, since the act of integrating the measured rate-of-turn to estimate heading acts by itself as a low-pass filter. However, in the HDR system a low-pass filter for smoothing the noisy gyro signal has a dramatic effect, since HDR acts on each individual gyro reading prior to the numeric integration. The low-pass filter is implemented in software

Heuristic reduction of gyro drift in vehicle tracking applications

$$\omega'_i = \frac{\omega'_{raw,i} T_i + \tau \omega'_{i-1}}{T_i + \tau} \quad (7)$$

where

- T_i – Sampling time
- ω' – Low-pass filtered rate of turn
- τ – Low-pass filter time constant

4.2.2 Turn Switch

Another enhancement to the basic HDR system, albeit a trivial one, is what we call the “Turn Switch.” When the vehicle takes a sharp turn, then the absolute value of the gyro-measured rate of turn, $|\omega'|$, spikes up to values that can be one or two orders of magnitude larger than any change in $|\omega'|$ due to drift. To prevent large rates of turn from affecting the Integrator I , a simple binary threshold can be used to distinguish between large $|\omega'|$ due to turning and small $|\omega'|$ that might be due to drift. If a reading of $|\omega'|$ exceeds that threshold, then the algorithm reduces the value of i_c to zero, for as long as $|\omega'|$ remains above the threshold. This measure effectively suspends the HDR algorithm, since HDR keeps incrementing I by zero until the condition is resolved. By doing so, HDR prevents itself from modifying I in response to signals that are caused by true turning, not by drift.

$$W_i = \begin{cases} 0 & \text{for } \omega_{i-1} < \Theta_w \\ 1 & \text{for } \omega_{i-1} \geq \Theta_w \end{cases} \quad (8)$$

where

- W_i – Turn switch
- Θ_w – Turn threshold

4.2.3 Repetition Attenuator

While sharp turns are easy to identify and handle by the Turn Switch, a challenge for the HDR algorithm is the difficulty of distinguishing between the effect of curving (e.g., driving along a slightly curved road) and drift. Curving is highly detrimental to the HDR algorithm, as illustrated in Figure 3. Figure 3 depicts an excerpt of a simulated drive that comprises three motion components: straight-line driving up to $t = 659$ sec, followed by slow curving at $2^\circ/\text{sec}$ from $t = 660$ to $t = 679$ sec, followed by straight-line driving. We

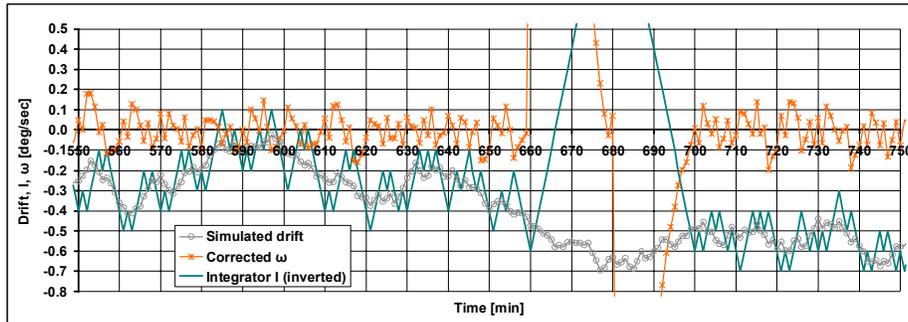


Figure 3: A simulated straight-line drive, interrupted by moderate curving at $2^\circ/\text{sec}$ from $t = 660$ sec to $t = 679$ sec. The curving causes a large spike in $-I$, which therefore no longer tracks drift. This introduces a large error in ω , namely, the negative peak between $t = 680$ sec and $t = 700$ sec.

chose a simulated drive for the discussion of Figure 3 because the simulation allows us to show the simulated and thus known drift, whereas in reality, of course, drift is not known. With $\omega = 2^\circ/\text{sec}$ during turning, its peak is not shown in Figure 3. Also, ω appears ragged because we simulated it with a very slight sinusoidal swerve that has a frequency of 0.1 Hz. Furthermore, since the plotted ω is the corrected ω (i.e., the output of the HDR control loop), it includes the oscillations of I (explained below). The light grey curve in Figure 3 is the simulated drift, ε_d , and the dark, thick curve is a plot of the output of the Integrator I , but with a negative sign. We plotted $-I$ instead of I because, as expected from the closed-loop HDR control system, $-I$ tracks the drift in steady state and during straight line driving, and that becomes obvious only when the two curves overlay each other as is the case when plotting $-I$.

If we assume for a moment that the Turn Switch threshold, as discussed in Section 4.2.2, is above $2^\circ/\text{sec}$, then the HDR algorithm cannot distinguish between curving and drift by means of the threshold. Consequently, HDR keeps incrementing $|-I|$ by i_c for the 20 seconds of curving. This results in a steeply increasing $|-I|$ so that $-I$ does not at all track drift, but rather chases ω . Once the curving ends at $t = 680$ sec, $|-I|$ peaks at $1.2^\circ/\text{sec}$ (not shown in Figure 3). This is a very large error in $|-I|$ and highly undesirable, since $-I$ is supposed to track drift. Moreover, since ω includes I , ω assumes a value of $-1.2^\circ/\text{sec}$ at $t = 680$ sec, instead of its nominal value of zero. With ω now negative, $-I$ chases it downward, thereby increasing ω . With $\omega_{true} = 0$, at $t \cong 700$, ω becomes zero again once $-I$ intersects the drift curve. From that point on $-I$ oscillates around the drift and steady state is reached. The episode, however, introduced an undesirably large error in ω .

To rectify this problem, we make use of the observation that straight-line motion, where $-I$ properly tracks drift, and curving, which misleads the HDR algorithm, can be distinguished by looking at the number of time intervals during which ω retains its sign. This is because at steady state, during straight line travel, ω tends to oscillate about zero as $-I$ oscillates about drift, as is evident in Figure 3. In contrast, during curving, ω retains its sign, because $-I$ chases ω .

With that distinguishing characteristic in mind, we can now implement the third enhancement for the HDR algorithm: the Repetition Attenuator, R . The task of R is to gradually reduce the value of the increment i_c , for every interval in which the sign of ω remained unchanged. For that, we first define a repetition counter, r . r is incremented by '1' for every iteration, in which the sign of ω_{i-1} remains unchanged from that of ω_{i-2} (note that since r_i is used to compute ω_i we have to go back two iterations, to ω_{i-2} , to see if the sign has changed).

$$r_i = \begin{cases} r_{i-1} + 1 & \text{for } \text{SIGN}(\omega_{i-1}) = \text{SIGN}(\omega_{i-2}) \\ 1 & \text{for } \text{SIGN}(\omega_{i-1}) \neq \text{SIGN}(\omega_{i-2}) \end{cases} \quad (9)$$

Then, the original increment i_c is reduced in some inverse proportion to r . In our system we created this inverse proportion by this function:

$$R_i = \frac{1 + c_1}{1 + c_1 r_i^{c_2}} \quad (10)$$

where

c_1 – Repetition Attenuator constant

c_2 – Repetition Attenuator power

Equation (10) allows the shaping of an attenuation curve by adjusting the parameters c_1 and c_2 . One such shape is depicted in Figure 4, where $c_1 = 0.16$ and $c_2 = 1.8$.

By multiplying i_c by R in every interval, i_c is effectively attenuated as a function of the repetition count r . The effect of this measure is that $|I|$ is incremented by the full,

Heuristic reduction of gyro drift in vehicle tracking applications

nominal value i_c only immediately after the sign of ω changes. If there is no change of sign for several intervals, I becomes almost stagnant. As a result, continuous curving cannot continue to affect I for more than a few intervals. This effect is illustrated in Figure 5. Here, thanks to the Repetition Attenuator enhancement, the area between the $-I$ curve and the drift curve is much smaller than the corresponding area in Figure 5. The Repetition Attenuator is highly effective, and c_1 and c_2 should be tuned carefully.

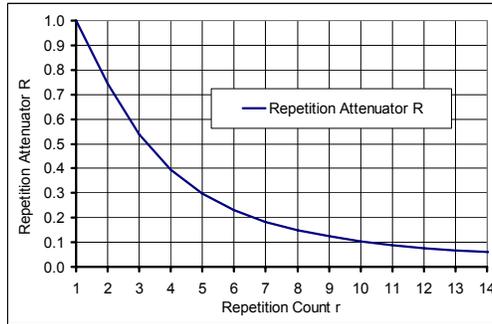


Figure 4: Plot of the Repetition Attenuator function of Eq. 10 for $c_1 = 0.16$ and $c_2 = 1.8$.

Discussion

It is generally undesirable for closed-loop control systems to generate oscillatory control signals. However, the oscillations of $-I$ about the actual but unknown drift are desirable in this application, for the following reason. I oscillates because the sign of the corrected ω oscillates about zero. As long as this is the case, the algorithm can assume that the vehicle is driving straight. Periods of time during which $-I$ is not oscillating are an indication that the vehicle is curving, as shown in Figure 3 and Figure 5. The repetition counter exploits this distinction by gradually reducing the increment i_c in prolonged periods of non-oscillation. The result is that the HDR algorithm does (almost) nothing during prolonged curving, but is active during straight-line travel.

We should also point out that the oscillations of the corrected ω are not harmful, because we are ultimately not interested in ω , but in heading, that is, the integral of ω with respect to time. This integral smoothes the oscillations of ω .

4.3 The Enhanced HDR Algorithm

The enhancements discussed in Section 4.2 can now be combined and substituted into the original Eqs. 4b and 5:

$$I_i = I_{i-1} - \text{SIGN}(\omega_{i-1})W_i R_i i_c \quad (11)$$

and

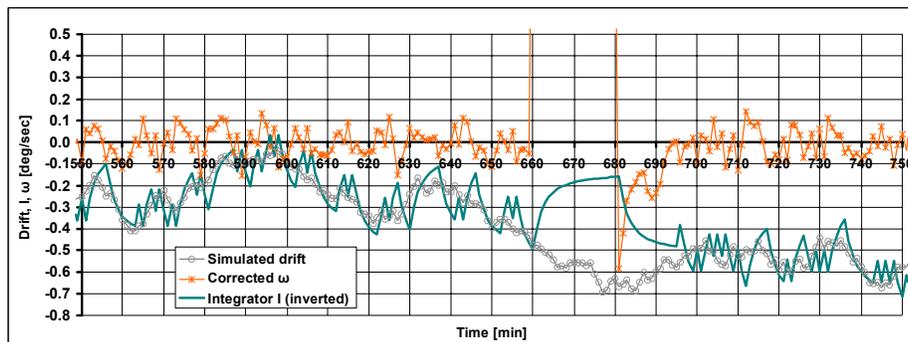


Figure 5: Same plot as in Figure 3, but now i_c is reduced adaptively with the Repetition Attenuator. The result is a much smaller error in ω .

$$\omega_i = \omega'_i + I_i \quad (12)$$

where ω'_i is the output of the low-pass filter implemented by Eq. 7.

The HDR-corrected rate of turn, ω_i , as computed by Eq. 12, still has one potential problem: The low-pass filter of Section 4.2.1, which was needed to smooth the noisy gyro signals, leaves ω_i with a small lag behind the raw ω'_i . If left this way, the vehicle's trajectory based on ω_i would be a slightly inaccurate representation of the true trajectory, as sharp turns become more drawn out. However, this problem can be effectively solved by reversing the effect of the low-pass filter. To this end, we invert Eq. 7:

$$\omega_{d,i} = \omega_i + \frac{\tau}{T_i}(\omega_i - \omega_{i-1}) \quad (13)$$

where

$\omega_{d,i}$ – The final, HDR-corrected and de-lagged rate of turn.

We can now compute the corrected heading ψ by rewriting Eq. (3)

$$\psi_i = \psi_0 + \psi_{i-1} + \omega_{d,i} T_i \quad (14)$$

where

ψ_0 – initial heading at the beginning of the drive

A block diagram of the complete HDR system is shown in Figure 6.

Due to space limitations, we omit here the discussion of some other enhancements, which provide additional, modest performance gains.

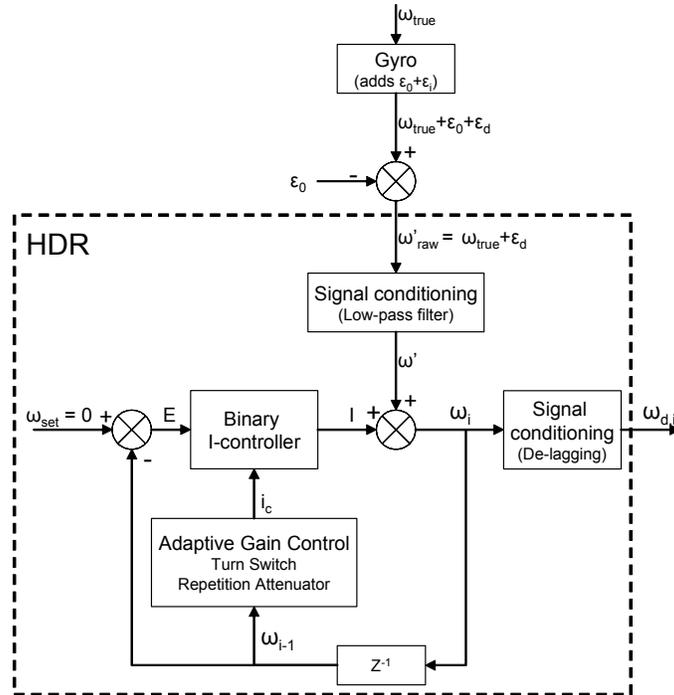


Figure 6: Final block diagram of the enhanced HDR system.

5 Experimental Results

In order to evaluate the effectiveness of HDR for vehicle tracking, we mounted a single-axis rate gyro on the floor of a small sports utility vehicle (SUV), a Subaru Forester. The gyro is the CRS03-04 made by Silicon Sensing [2008], shown in Figure 7. Table I lists some of the key specification for the CRS03-04. In all experiments we sampled the gyro data at a rate of 10 Hz, while driving. Simultaneously we logged GPS data to serve as ground truth. Since the GPS data was sampled at 1 Hz, we down-sampled the Gyro data to 1 Hz by averaging every 10 samples



Figure 7: The Silicon Sensing CRS03-04 single-axis rate gyro.

$$\omega_i^{1Hz} = \frac{1}{10} \sum_{j=1}^{10} \omega_{10i+j}^{10Hz} \quad \text{for } i = 1, 2 \dots n \quad (14)$$

Consequently, all experimental results in this paper were obtained from computations performed on the basis of a sampling time of $T = 1$ sec.

5.1 A Typical Drive Experiment

The nominal route of a typical drive experiment with HDR is depicted in Figure 8. This drive, called the “Mixed Drive,” comprised of 8 km of rural roads (from ‘Start’ to ‘A’), 11.5 km of highway (from ‘A’ to ‘B’), and 3.75 km of city streets (from ‘B’ to ‘Stop’). The route was purposefully designed to include stretches of highway that are curving, not just straight segments. The total route was 23.6 km long and took ~20 minutes to drive.

Figure 9 shows a plot of heading versus time for the Mixed Drive experiment. The thick blue curve represents ground truth heading computed from GPS data. Because of the variance of the GPS updates, heading estimation at slow speeds was often very noisy. Whenever that was the case, we smoothed the noisy GPS data manually. The dotted, red curve shows heading as estimated from uncorrected gyro data, while the thin green curve shows heading estimated with HDR-corrected data.

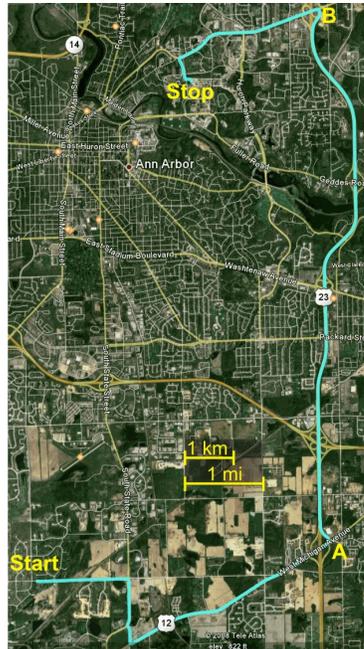


Figure 8: The nominal route for the ‘Mixed Drive’ experiment.

Table I: Key specifications for the Silicon Sensing CRS03-04.

Size	29×29×18 mm
Bandwidth	10 Hz
Dynamic range	±200 deg/sec
Bias drift	<1.1 deg/min (66 deg/hr)
Cost	350 USD

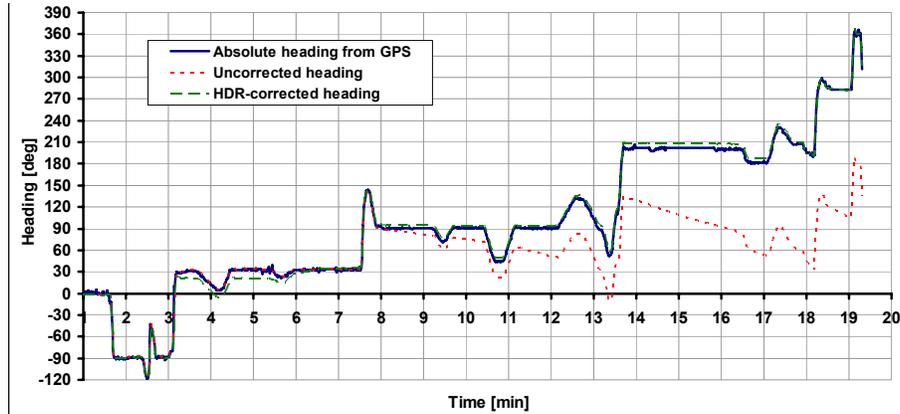


Figure 9: A plot of heading data from the Mixed Drive experiment of Section 5.1. The thick blue curve shows ground truth, derived from GPS data. The dotted red curve shows heading estimated from gyro data, without HDR correction. The dashed green curve shows heading estimated after applying the HDR algorithm.

Note how closely the HDR-corrected curve follows the ground truth data. Quantitative results for this drive and others are given in the following section.

5.2 Multiple Drive Experiments

It is quite easy to adjust the tunable parameters of the HDR algorithm so as to provide excellent performance for a single, specific drive experiment, in post-processing. In practice, however, parameters tuned for optimal performance in one experiment will almost certainly not produce optimal performance in another one. The only practical way for tuning parameters is to collect data from a large number of experiments and observe how a single set of parameters performs when applied to all available data sets. Parameters tuned this way will not provide optimal performance for any individual experimental drive, but they are far more likely to produce acceptable results with any future drive in a production vehicle tracking system, where re-tuning the parameters is out of the question. One condition for this to be true is that the so-called “training” data sets be based on a wide range of different drive conditions.

With these considerations in mind, we performed nine different drive experiments under a wide range of driving conditions. Specifically, the drives included highways, rural roads, and city streets, and they varied between 13 and 52 minutes in duration and between 11 and 91 km in distance. Each experiment started by measuring the static bias drift, ε_0 , while the vehicle was standing still for about 20 seconds. In order to express the results of these experiments quantitatively, we use two metrics: the *Average Heading Error*, E_ψ , and the *Normalized Average Heading Error*, NE_ψ , described next.

5.2.1 The Average Heading Error (E_ψ)

This metrics is useful for judging the accuracy achieved in a single drive. In order to compute E_ψ we first computed the momentary heading error by comparing the gyro-derived estimated heading with the GPS-derived ground truth heading in each sampling interval (i.e., once every second). Then, the averaging of all momentary heading errors yields the Average Heading Error:

Heuristic reduction of gyro drift in vehicle tracking applications

$$E_{\Psi} = \frac{1}{n} \sum_{i=1}^n |\psi_{i,Gyro} - \psi_{i,GPS}| \quad (15)$$

where

E_{Ψ} – Average Heading Error in degrees.

i – Index for the interval number within each drive.

n – Total number of samples in the drive (the effective sampling frequency was 1 Hz).

$\psi_{i,Gyro}$ – Gyro-derived heading in interval i , in degrees.

$\psi_{i,GPS}$ – GPS-derived heading in interval i , in degrees.

5.2.2 The Normalized Average Heading Error (NE_{Ψ})

Since heading errors due to drift can increase without bound, one can expect to see only small momentary heading errors shortly after starting a drive, and increasingly larger momentary heading errors later into the drive. Consequently, one can expect the Average Heading Error to be smaller for short-duration drives and larger for long-duration drives. In order to be able to compare the performance between short and long drives, we define a second metrics, the *Normalized Average Heading Error per minute of travel time*, NE_{Ψ} , which is useful for comparing the accuracy achieved in multiple drives of different durations. NE_{Ψ} is measured in degrees per minute.

$$NE_{\Psi} = \frac{E_{\Psi}}{T_d} \quad (16)$$

where

T_d – duration of the drive in minutes.

Table II lists relevant details and results for all nine drives. A graphical representation of these results is shown in Figure 10. For each drive we specify both E_{Ψ} and NE_{Ψ} .

We should emphasize that a single set of parameters was used for all drives, and all results of Table II and Figure 10 were obtained with that fixed set of parameters. These parameters and their values are listed in Table III.

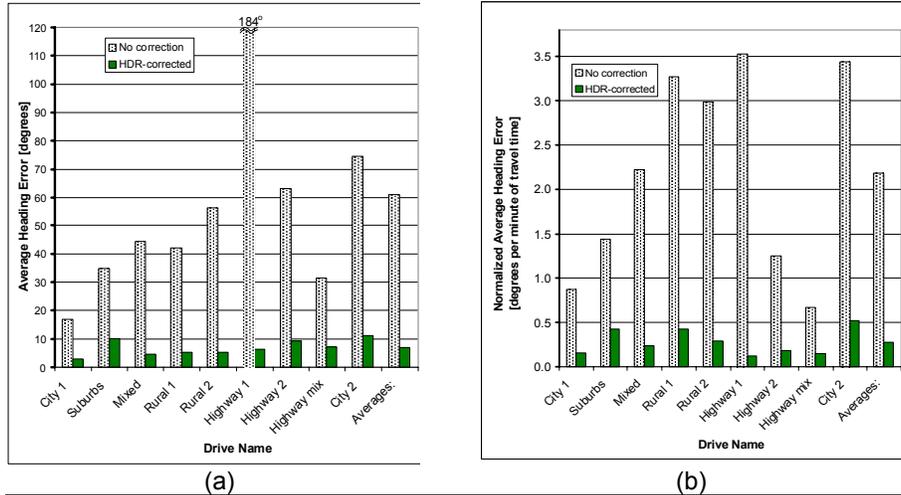


Figure 10: Graphical representation of the results of nine drive experiments with and without HDR correction. (a) Average Heading Errors E_{Ψ} ; (b) Normalized Average Heading Errors, NE_{Ψ} .

Table II: Description and results for the nine drive experiments.

Experiment	Duration	Distance	Description	Average Heading Errors		
				Without HDR	With HDR	Relative HDR Improvement
City 1	19.4 min	10.8 km	Mostly straight streets, few turns	$E_{\psi}: 17^{\circ}$ $NE_{\psi}: 0.87^{\circ}/min$	$E_{\psi}: 3.1^{\circ}$ $NE_{\psi}: 0.16^{\circ}/min$	5.4×
Suburbs	24.3 min	17.0 km	Mostly curving streets in the suburbs	$E_{\psi}: 35^{\circ}$ $NE_{\psi}: 1.4^{\circ}/min$	$E_{\psi}: 10.4^{\circ}$ $NE_{\psi}: 0.43^{\circ}/min$	3.4×
Mixed*	20.0 min	23.6 km	Mix of city streets, rural roads, and highway	$E_{\psi}: 44^{\circ}$ $NE_{\psi}: 2.2^{\circ}/min$	$E_{\psi}: 4.7^{\circ}$ $NE_{\psi}: 0.24^{\circ}/min$	9.4×
Rural 1	12.9 min	16.1 km	Mostly straight streets, few turns	$E_{\psi}: 42^{\circ}$ $NE_{\psi}: 3.3^{\circ}/min$	$E_{\psi}: 5.5^{\circ}$ $NE_{\psi}: 0.42^{\circ}/min$	7.7×
Rural 2	18.9 min	15.5 km	Mostly straight rural roads, some curving suburban streets	$E_{\psi}: 56^{\circ}$ $NE_{\psi}: 3.0^{\circ}/min$	$E_{\psi}: 5.5^{\circ}$ $NE_{\psi}: 0.29^{\circ}/min$	10.2×
Highway 1	52.1 min	91.4 km	90% highway, mostly straight	$E_{\psi}: 184^{\circ}$ $NE_{\psi}: 3.5^{\circ}/min$	$E_{\psi}: 6.4^{\circ}$ $NE_{\psi}: 0.12^{\circ}/min$	29×
Highway 2	50.4 min	80.6 km	90% highway, some curving	$E_{\psi}: 63^{\circ}$ $NE_{\psi}: 1.3^{\circ}/min$	$E_{\psi}: 9.3^{\circ}$ $NE_{\psi}: 0.19^{\circ}/min$	6.8×
Highway Mix	47.6 min	81.2 km	70% highway with some curving, 30% rural, mostly straight	$E_{\psi}: 32^{\circ}$ $NE_{\psi}: 0.66^{\circ}/min$	$E_{\psi}: 7.1^{\circ}$ $NE_{\psi}: 0.15^{\circ}/min$	4.4×
City 2	21.6 min	11.9 km	100% downtown streets, over 40 90-degree turns	$E_{\psi}: 74^{\circ}$ $NE_{\psi}: 3.4^{\circ}/min$	$E_{\psi}: 11.3^{\circ}$ $NE_{\psi}: 0.52^{\circ}/min$	6.7×
Total for all nine drives:	4.45 hrs	348 km	Average of all nine drives:	$\bar{E}_{\psi}: 61^{\circ}$ $\bar{NE}_{\psi}: 2.2^{\circ}/min$	$\bar{E}_{\psi}: 7.0^{\circ}$ $\bar{NE}_{\psi}: 0.28^{\circ}/min$	9.2× $\bar{\sigma} = 7.75$

*Note: This is the Mixed Drive experiment that was described in more detail in Section 5.1.

Discussion

Several issues related to the cited results deserve further discussion:

1. It is quite apparent that some of the uncorrected heading errors are substantially larger in some drives than in others. In particular, the Highway 1 drive has an uncorrected error of 184° , much more than what one might expect due to drift alone. To understand these large differences, one should consider the following additional error sources, each of which can cause errors many times larger than bias drift alone.
 - a. Some of the drives were performed at near-freezing outdoor temperatures. The static bias drift was measured shortly after the test vehicle was driven out of a garage, where the interior temperature was on the order of $5\text{-}10^{\circ}\text{C}$ ($41\text{-}50^{\circ}\text{F}$). Then, during the drive, the interior temperature of the vehicle was increased to room temperature, for driver comfort. This large change in temperature has a profound effect on drift, which may be much larger than the manufacturer-

Table III: A single set of parameters was used for all drives.

Parameter	Symbol	Value	Units
Nominal increment	i_c	0.0422	$^{\circ}/\text{sec}$
Low-pass filter time const.	τ	0.75	sec
Turn threshold	θ_w	0.564	$^{\circ}/\text{sec}$
Rep. attenuator constant	c_1	0.16	none
Rep. attenuator power	c_2	1.80	none

specified nominal bias drift rate.

- b. Even small changes in the way the static bias drift is measured before a drive can have profound effects on heading errors, especially on long drives. For example, in the Highway Mix drive, we measured a static bias drift of $\varepsilon_{0,1} = 1.432^\circ/\text{sec}$ and the average heading error, according to Table II, was $E_\psi = 32^\circ$. Had we measured a slightly different static bias drift, say, $\varepsilon_{0,2} = 1.446^\circ/\text{sec}$ (which differs from $\varepsilon_{0,1}$ by just 1%), then that would have yielded $E_\psi = 11^\circ$ in the same drive. Variations on the order of 1% or larger in measuring ε_0 are indeed possible. For example:
 - when the static bias drift test is done for less than the recommended amount of time;
 - due to the gyro not being perfectly level during the static bias test (MEMS gyros are sensitive to acceleration, e.g., gravitation); or
 - due to changes in temperature during the static bias test.
2. The Highway 1 Drive has an exceptionally strong Relative HDR Improvement score of 29-fold. That is almost three times higher than the next-best score of 10.2-fold for the Rural 2 Drive. This possible exception inflates the Average Relative HDR Improvement to 9.1-fold (bottom right cell in Table II). If we exclude the Highway 1 Drive from the overall average, then that revised Average Relative HDR Improvement would be reduced from 9.2-fold with a standard deviation of $\hat{\sigma} = 7.75$ to 6.6-fold with a standard deviation of $\hat{\sigma} = 2.35$. However, we do not see any compelling reason for actually excluding this drive. The uncorrected heading error, while large, is not much greater than that of other drives when normalized (i.e., divided by the duration of the drive), as is apparent from Figure 10b.

5.3 Test Set

The tunable parameters of the HDR algorithm were painstakingly hand-tuned to minimize the bottom line in Table II, and specifically the Average of all Normalized Average Heading Errors ($\overline{NE}_\psi = 0.28^\circ$). This is particularly apparent in Table III, which shows that some of the parameters were tuned with the precision of three significant digits. This raises the question of how well the algorithm would perform with additional drives, for which the parameters were not specifically tuned.

To answer this question, we performed another set of five drives and applied the HDR algorithm to those drives with the exact same parameter values as those of Table IV. The results are summarized in Table IV and Figure 11. Borrowing terminology from the field of Neural Networks, we call the set of drives used for tuning the parameters (i.e., the nine drives listed in Table II) the “Training Set,” and the set of drives, to which the fixed parameters of Table III were applied (i.e., the drives listed in Table IV), the “Test Set.” In the Test Set drives HDR provides a 6.9-fold improvement ($\hat{\sigma} = 3.44$) over uncorrected heading estimates.

6 Conclusions

In this paper we proposed the HDR method for reducing errors due to gyro drift in vehicle tracking applications. The basic heuristic is that much driving is done along reasonably straight lines. Whenever that is the case, the closed-loop control approach in our system leads the controller output, $-I$, to track the drift. Subtracting I from the gyro-measured rate of turn data then effectively removes the estimated drift.

Table IV: Description and results for the five drives of the Test Set.

Experiment	Duration	Distance	Description	Average Heading Errors		
				Without HDR	With HDR	Relative HDR Improvement
City 3	22.2 min	13.3 km	Mostly straight streets, few turns	E_{ψ} : 45° NE_{ψ} : 2.1°/min	E_{ψ} : 10.6° NE_{ψ} : 0.48°/min	4.3×
Mixed 2	19.2 min	14.2 km	Mix of city streets, rural roads, and highway	E_{ψ} : 35° NE_{ψ} : 1.8°/min	E_{ψ} : 10.5° NE_{ψ} : 0.55°/min	3.3×
Rural 3	16.6 min	13.3 km	Mix of city streets and rural roads	E_{ψ} : 44° NE_{ψ} : 2.7°/min	E_{ψ} : 3.8° NE_{ψ} : 0.23°/min	11.6×
Highway 3	36.9 min	62.8 km	Highway with many curving segments	E_{ψ} : 192° NE_{ψ} : 5.2°/min	E_{ψ} : 21.1° NE_{ψ} : 0.57°/min	9.1×
Highway 4	43.2 min	78.0 km	Highway with many curving segments	E_{ψ} : 142° NE_{ψ} : 3.3°/min	E_{ψ} : 22.5° NE_{ψ} : 0.52°/min	6.3×
Total for all five Test Set drives:	1.58 hrs	104 km	Average of all five Test Set drives:	\bar{E}_{ψ} : 92° \bar{NE}_{ψ} : 2.9°/min	\bar{E}_{ψ} : 13.7° \bar{NE}_{ψ} : 0.46°/min	6.9× ($\sigma = 3.44$)

Although in simulations HDR virtually eliminates the effects of drift (in real runs one cannot measure drift, thus making it impossible to prove or disprove this claim for real runs) some errors remain, due to two reasons:

1. The basic HDR algorithm cannot distinguish well between drift and actual curving motion. To reduce this undesirable behavior, we introduced enhancements: the Low-pass Filter, the Turn Switch, and the Repetition Attenuator. With these enhancements the HDR algorithm effectively suspends its operation for as long as the actual curv-

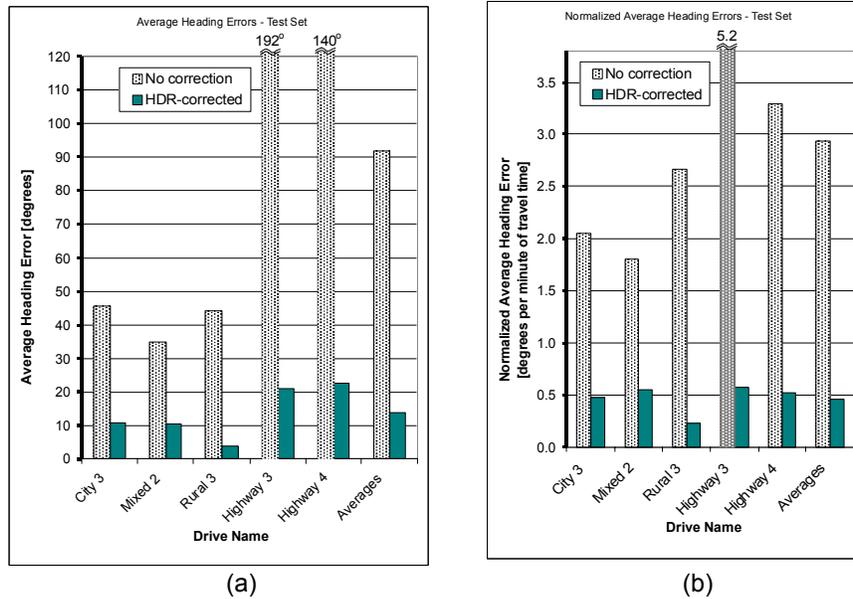


Figure 11: Graphical representation of the results of the five *Test Set* drive experiments with and without HDR correction. (a) Average Heading Errors E_{ψ} ; (b) Normalized Average Heading Errors, NE_{ψ} .

Heuristic reduction of gyro drift in vehicle tracking applications

ing motion continues. Then, when straight line motion resumes, HDR corrections resume and $-I$ resumes its tracking of drift. Heading errors due to drift while HDR was suspended cannot be recovered and they are the main contributor to errors in the HDR system.

2. Even during straight-line motion, there is always a residual offset between $-I$ and drift because drift never reaches steady state. This residual offset incurs heading errors that cannot be recovered.

Despite the errors that HDR introduces, its corrective effect outweighs its errors by a wide margin, as is apparent from the experimental results. For the gyro used in our vehicle tracking system, HDR reduced heading errors by a factor of 9.2 for the Training Set and a factor of 6.9 in the Test Set.

Other advantages of the HDR method are:

- The HDR method can be implemented in software only, in as little as 20 or so lines of code, and, besides the gyro, no additional sensors or data are needed.
- Another, potentially major advantage that we did not address in this paper is the fact that HDR can eliminate the need for measuring static bias drift for 20 seconds before each drive. In informal testing we found that just one second of measured static bias drift, e.g., the time between turning the ignition key and starting to drive, is sufficient for an HDR system to produce very usable results with errors just slightly larger than those reported in this paper.

Acknowledgements

This work was supported by the U.S. Department of Energy under Award No. DE FG52 2004NA25587.

7 References

- Basnayake, C., Mezentsev, O., Lachapelle, G., and Cannon, M.E. (2005) "An HSGPS, inertial and map-matching integrated portable vehicular navigation system for uninterrupted real-time vehicular navigation." *International Journal of Vehicle Information and Communication Systems*, Vol. 1, No.1/2 pp. 131–151.
- Cavallo, F. Sabatini, A.M., and Genovese, V. (2005) "A step toward GPS/INS personal navigation systems: real-time assessment of gait by foot inertial sensing." *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2-6 Aug., pp. 1187–1191
- Chen, X., 2004, "Modeling Temperature Drift of FOG by Improved BP Algorithm and by Gauss-Newton Algorithm." *Lecture Notes in Computer Science – Springer Berlin / Heidelberg*, ISBN 978-3-540-22843-1, pp. 805–812.
- Cho, S.Y. Lee, K.W. Park, C.G., and Lee, J.G., 2003, "A Personal Navigation System Using Low-Cost MEMS/GPS/Fluxgate." *Proceedings of the 59th Institute of Navigation (ION) Annual Meeting*, June 23-25, Albuquerque, NM.
- D.A. Grejner-Brzezinska, C. Toth, S. Moafipoor, Y. Jwa, J. Kwon, 2006, "Multi-sensor personal navigator supported by human motion dynamics model." *3rd IAG / 12th FIG Symposium*, Baden, Austria, May.
- IEEE Standards, 1996, "IEEE Standard Specification Format Guide and Test Procedure for Single-Axis Interferometric Fiber Optic Gyros."
- Mohinder S. Grewal, Lawrence R. Weill, Angus P. Andrews, 2002, "Global Positioning Systems, Inertial Navigation, and Integration" Copyright 2001 John Wiley & Sons, Inc.

Johann Borenstein and Lauro Ojeda

Pandit, S.M. and Zhang, W., 1986, "Modeling Random Gyro Drift Rate by Data Dependent Systems." *IEEE Transactions on Aerospace and Electronic Systems*, Vol. AES-22, No. 4, July, pp. 455-460.

Silicon Sensing (2008) <http://www.siliconsensing.com/CRS03packaged>, last accessed June 2008.