Lane-Keeping Safety and Time Correlation of Navigation Errors

Olivier N. Kigotho and Jason H. Rife, *Tufts University*
Hadi Wassaf, *USDOT Volpe Center*

BIOGRAPHY

Olivier Ng'weno Kigotho is a staff researcher in the Department of Mechanical Engineering at Tufts University in Medford, Massachusetts. He supports the Automated Systems and Robotics (ASAR) Laboratory, which applies theory and experiment to characterize integrity of autonomous vehicle systems. He received his B.S. in Mechanical Engineering at Tufts University.

Jason Rife is Professor and Chair of the Department of Mechanical Engineering at Tufts University in Medford, Massachusetts, where he directs the ASAR Laboratory. He received his B.S. in Mechanical and Aerospace Engineering from Cornell University and his M.S. and Ph.D. in Mechanical Engineering from Stanford University.

Dr. Hadi Wassaf is an Electrical Engineer in the Aviation Weather and Positioning, Navigation, and Timing (PNT) Applications Division of the U.S. DOT Volpe National Transportation Systems Center (Volpe Center). Dr. Wassaf leads the PNT for automated systems work at the Volpe Center in support of the Office of the Assistant Secretary of Transportation for Research and Technology (OST-R) PNT directorate and the Intelligent Transportation Systems Joint Program Office (ITS JPO). Dr. Wassaf holds a BS, MS, and PhD in Electrical Engineering from the University of Massachusetts Lowell.

ABSTRACT

Recent work has constructed alert limits for autonomous vehicle lane-keeping applications, where safety criteria were based on the sensor readings of the navigation system. One limitation of this prior work is that measurement errors were compared directly to lane and vehicle geometry, with the assumption that measurement errors would map directly to errors in vehicle position or attitude. In fact, measurement errors pass through a dynamic system that includes many processes that transform the error, including natural processes that correlate sensor measurements over time and closed-loop vehicle dynamics. These processes reshape error distributions, such that the true position and heading error distributions induced by sensor noise have different forms than the original measurement error distributions. This paper explores these effects when the measurement error is time correlated. To this end, we abstract the dynamic system using a unicycle model of vehicle dynamics, a representative controller, and a first-order Gauss-Markov process to represent measurement correlation. Implementing this model in a Monte Carlo simulation, we investigate the case of a vehicle moving on a curved road when the time constant for measurement correlation is moderately long (10 seconds). Simulations show that the effects of heading-measurement error are mitigated by the dynamic system. Meanwhile, the two-dimensional position error distribution rotates relative to the road, an effect which makes position errors more difficult to bound.

INTRODUCTION

This paper focuses on safety for automobile navigation systems, with a particular emphasis on automated lane-keeping. We consider how time-correlated errors impact lane-keeping safety requirements. Our analysis considers two processes that introduce time-dynamics into the vehicle state error: namely measurement dynamics and vehicle dynamics.

Our analysis builds on prior research that have analyzed the mapping between navigation system error and safety requirements. This paper extends the analyses of Reid [1] and Kigotho [2], which derived sensor-error bounds called

alert limits that ensure safety for lane-keeping applications on straight and curved roads. For lane-keeping on straight roads, the largest tolerable lateral-positioning errors are dictated by lane and vehicle width. For curved roads, errors must be modeled in two dimensions, with the shape of the road coupling the tolerable error levels in the lateral and longitudinal direction. Both [1], which used rectangular alert limits, and [2], which used elliptical alert limits, showed that road curvature reduces the allowable lateral error, with less tolerance for lateral uncertainty when longitudinal uncertainty increases. In both [1] and [2], sensor measurement errors were considered only in an instantaneous sense, with no consideration for the evolution of the sensor noise and the vehicle trajectory in time.

For safety-critical applications, the alert limit should consider time correlation of state errors. The alert limit would be used in a screening process that can introduce an alarm given too high a probability of a large measurement error, where the measurement error is considered *large* if it falls outside the alert limit. This screening process ensures integrity, but at the expense of availability, which is the fraction of the time that the system meets its safety requirements and is therefore usable. Typically, error probabilities are assessed by defining a *protection level*, which is a confidence bound derived from an error model at a specified probability that reflects an acceptable level of integrity risk. To ensure safety, errors are commonly modeled as *overbounds*, models that conservatively represent the far distribution tails [3]–[6]. Overbounds are usually modeled as static processes, including conservative assumptions to cover nonstationary signals, characterized by distributions that change in time [7]. But these models do not generally cover cases of time-correlated errors, where bad error states at one time may imply even worse error states later. Time-correlation issues are known to impact navigation systems, however. For instance, GNSS multipath positioning errors, introduced when signals reflect off surfaces like terrain or buildings, may be correlated over times of several seconds or even minutes, depending on the relative motion between the GNSS receiver and the reflecting surface [8], [9].

To account for time-correlation effects, our approach is to model dynamic processes including measurement and vehicle dynamics. Measurement dynamics map a random signal to a navigation system error. Vehicle dynamics map the navigation system error to a vehicle state error, which can then be compared to the safety requirement (e.g. an alert limit). Both processes can result in time correlation.

Our primary contribution in this paper is to identify trends that characterize how these dynamic processes influence the state errors that must be compared to alert limits to ensure safety. Using simulations, we show that vehicle dynamics can mitigate heading-measurement errors, meaning that the heading state-error distribution is narrower (less uncertain) than the measurement-error distribution. Also, our simulations indicate that system dynamics correlate along-track errors and cross-track errors for vehicles traveling on curved roadways. Both effects have significant impact for evaluating the integrity risk associated with navigation system errors.

The rest of the paper is organized as follows. In the next section we discuss the low-order model of system dynamics that we use to map measurement errors to position and attitude errors. We then detail the Monte Carlo simulations used to evaluate the dynamic model. The final sections of the paper present our results and discuss their significance.

SYSTEM DYNAMICS MODEL

In this section, our goal is to define a representative, low-order model of the dynamic system that transforms navigation system errors into vehicle state errors. In this paper we use the term *state* to collectively refer to the variables describing a vehicle's actual position and attitude. The term *state error* refers to deviations of the actual vehicle position and attitude from the intended trajectory.

A detailed system dynamic model would be dependent on many factors including the specific vehicle model and the operating environment. Relevant dynamic processes include environmental effects that correlate sensor noise over time (e.g. multipath for GPS [8], [9]), sensor hardware dynamics (e.g. the rotation of a spinning lidar system), sensor fusion that can introduce filtering lags, and dynamics associated with control algorithms, actuators, and the vehicle platform. A qualitative block diagram that illustrates these processes is shown in Fig. 1. Note that the feedback loop results in the vehicle state, which must be compared to safety requirements to assess risk. These requirements are represented as a block on the right side of Fig. 1.

Although a detailed dynamic model might be required to obtain a highly accurate prediction of performance for a specific vehicle and operating environment, we assert that a simplified model can provide descriptive information about the behavior of the breadth of different autonomous vehicle systems. For instance, whether the vehicle controller is designed using machine-learning methods or analytical methods, the controller will seek to reduce state error; in that sense, a canonical controller, such as a proportional-derivative controller, can be used to represent the typical behavior of such systems. Using such approximations of vehicle system dynamics, this section proposes a low-order model that we subsequently use for analysis.

Our simplified vehicle model consists of two processes: measurement and vehicle dynamics. In effect, we transform Fig. 1 into a simplified form, illustrated in Fig. 2. The simplified block diagram focuses purely on the response of the vehicle to sensor noise. As such, we do not consider other forcing inputs such as driver control commands or external disturbance forces on the vehicle (e.g. wind, potholes). We focus instead on modeling dynamics related to the measurements (e.g. sensor model, noise processes, and sensor fusion) together in one block and dynamics associated with the vehicle itself (control and vehicle dynamics including actuation) together in a second block. Using conventional block diagram manipulation [10],the feedback loop can be embedded within the vehicle dynamics block. The vehicle dynamics result in the state perturbations due to sensor noise; these state-errors can be compared to the safety requirement to assess risk.

Though greatly simplified compared to an actual system, the system in Fig. 2 captures important processes that have been left out of earlier research related to mapping sensor errors to safety requirement for lane-keeping [1],[2]. Prior work has mapped measurement errors directly to safety requirements, as shown in Fig. 3, without considering measurement or vehicle dynamics.

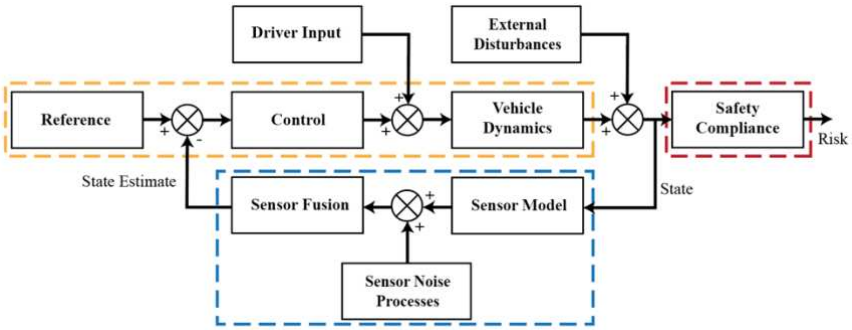

Fig. 1. Qualitative block diagram for system dynamics governing lane-keeping safety

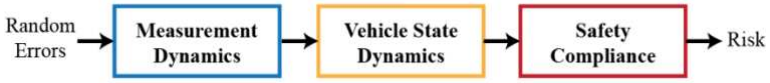

Fig. 2. Abstracted block diagram used in our analysis

Fig. 3. Abstracted block diagram used in prior analyses [1],[2]

The remainder of this section describes the vehicle and measurement dynamic models used to represent the related blocks in Fig. 2.

Vehicle Dynamics

To capture the dynamics of an autonomous vehicle, we assume a kinematic model, sometimes called a unicycle model, where the heading θ and position vector $\mathbf{x}$ are controlled by setting velocity v and angular velocity ω. The position

vector is described in Cartesian coordinates (indicated below by the leading superscript c). The time-evolution of the vehicle states are governed by

$$\dot{\mathbf{x}} = {}^{c}\begin{bmatrix}\dot{x}_c\\ \dot{y}_c\end{bmatrix} = {}^{c}\begin{bmatrix}v\cos\theta\\ v\sin\theta\end{bmatrix} \quad \text{and} \quad \dot{\theta} = \omega \tag{1}$$

For analytical purposes, it is also useful to represent the vehicle states in path coordinates, where the coordinates are expressed relative to a nominal trajectory. The nominal trajectory is defined in terms of a position vector $\mathbf{x}_{ref}$ and a scalar heading θ_{ref}, which are both functions of time. For example, the position state associated with the vehicle center is the vector from the reference at point O_p to the actual vehicle center at point O_v, as shown below in Fig. 4. In path coordinates, the position vector is described using a basis that instantaneously aligns with the reference trajectory at point O_p, with the coordinate x_p describing the direction locally tangent to the path and the coordinate y_p describing the direction locally normal to the path. These measures of longitudinal and lateral (or along-track and cross-track) error can be combined together into the vector $\mathbf{x}_p$ which represents the state error.

$$\mathbf{x}_p = \mathbf{x} - \mathbf{x}_{\text{ref}} \tag{2}$$

By using Euler's rule for vector differentiation, the unicycle position dynamics of (1) can be expressed in path coordinates as

$$\dot{\mathbf{x}}_p = \frac{{}^{c}\mathrm{d}}{\mathrm{d}t}(\mathbf{x} - \mathbf{x}_{\text{ref}}) - \boldsymbol{\omega}_{\text{ref}} \times (\mathbf{x} - \mathbf{x}_{\text{ref}}) = {}^{p}\begin{bmatrix} v\cos(\theta - \theta_{\text{ref}}) - v_{\text{ref}} + \dot{\theta}_{\text{ref}}\, y_p \\ v\sin(\theta - \theta_{\text{ref}}) - \dot{\theta}_{\text{ref}} x_p \end{bmatrix} \tag{3}$$

The cross product is formed here using an angular velocity vector $\boldsymbol{\omega}_{\text{ref}}$, which is equal to the reference angular rate $\omega_{\text{ref}} = \dot{\theta}_{\text{ref}}$ multiplied by the upward-pointing unit vector.

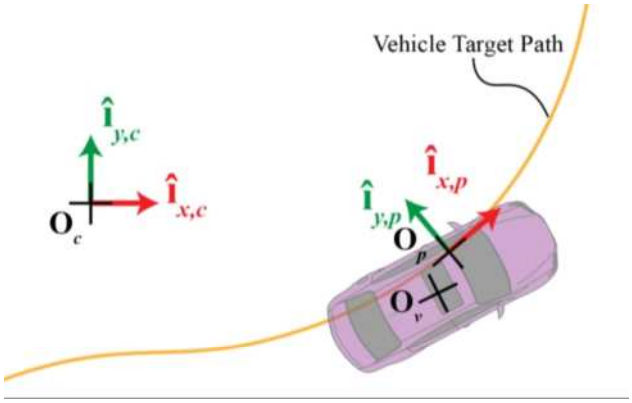

Fig. 4 Coordinate systems

The baseline vehicle trajectory is achieved by controlling the vehicle speed v and turn rate ω. Ideally, for a feasible trajectory and in the absence of noise, v and ω can be specified precisely to follow the desired trajectory $\mathbf{x}_{ref}$. We refer to these idealized control inputs as the open-loop (or *feedforward*) control commands. To ensure consistency with the dynamic model (1), we define the reference trajectory by first choosing the reference velocity v_{ref} and angular velocity ω_{ref} as functions of time. Appealing to (1), the reference angle $\theta_{\text{ref}}(t)$ can be obtained trivially by integrating $\omega_{\text{ref}}(t)$, and the reference position $\mathbf{x}_{ref}(t)$ can be obtained by integrating

$$\dot{\mathbf{x}}_{ref} = {}^{c}\begin{bmatrix}\dot{x}_{\text{ref}}\\ \dot{y}_{\text{ref}}\end{bmatrix} = {}^{c}\begin{bmatrix}v_{\text{ref}}\cos\theta_{\text{ref}}\\ v_{\text{ref}}\sin\theta_{\text{ref}}\end{bmatrix} \tag{4}$$

Feedforward control is insufficient, however, to reject disturbances such as sensor noise or external forces that push the vehicle away from the reference trajectory. Feedback control is needed to observe and reject disturbances.

The goal of the feedback controller is to keep the vehicle on its desired trajectory. In a modern autonomous vehicle system, the feedback controller may be a complex architecture involving hybrid control and/or elements of machine learning. Despite the wide range of design implementations, the objective of lane-keeping controllers is always more-or-less the same: to keep the vehicle on a target trajectory away from the lane boundaries. To this end, we represent this basic behavior through a linear control model. The linear model captures the control objective in a manner that is computationally simple, to make our analysis as tractable as possible.

Our model controller keeps the vehicle on path by applying proportional gains k_x, $k_y v_{\text{ref}}$, and $k_\theta v_{\text{ref}}$ to eliminate observable state errors, which include $\mathbf{x}_p$ and the relative heading $\theta - \theta_{\text{ref}}$. The control gains are scaled by the target velocity v_{ref} to help keep the vehicle from overshooting its path at high speeds. Including the feedforward and feedback terms together, the linear control law is

$$\begin{aligned} v &= v_{\text{ref}} - k_x x_p \\ \omega &= \omega_{\text{ref}} - k_y v_{\text{ref}}(y_p) - k_\theta v_{\text{ref}}(\theta - \theta_{\text{ref}}) \end{aligned} \tag{5}$$

Note that the controller can only respond to error signals that can be observed through sensing. Sensor noise, particularly when time-correlated, can mask true errors and result in a nonideal control signal. To account for this, we introduce the scalar errors ε_x, ε_v, ε_θ, ε_y, and $\varepsilon_\omega$to account for sensor noise. The modified control law has the form

$$\begin{aligned} v &= (v_{\text{ref}} + \varepsilon_v) - k_x(x_p + \varepsilon_x) \\ \omega &= (\dot{\theta}_{\text{ref}} + \varepsilon_\omega) - k_y v_{\text{ref}}(y_p + \varepsilon_y) - k_\theta v_{\text{ref}}((\theta + \varepsilon_\theta) - \theta_{\text{ref}}) \end{aligned} \tag{6}$$

Measurement Dynamics

A typical model for stochastic errors used to model navigation sensor data is the Gauss-Markov process (GMP) [11]. Since we do not know for certain what distribution of uncertainty the sensors would produce, a starting point would be to model the errors using target output Gaussian distributions for ε_x, ε_v, ε_θ, ε_y, and ε_ω. To model time correlation in these errors, we use a first order GMP. The GMP creates time-correlated errors by passing a random time-uncorrelated signal through a low pass filter, which abstractly represents natural processes in the environment that correlate the signal arriving at the sensors. The first-order system has the following form.

$$\dot{\varepsilon}_i + \frac{1}{\tau}\varepsilon_i = \nu \tag{7}$$

Here ν is the uncorrelated noise driving the measurement error, and ε_i is each of the errors ε_x, ε_v, ε_θ, ε_y, and ε_ω.

SIMULATION

In this section we describe how we implemented the simulation. Our objective is to transform a distribution of sensor errors into a distribution of system states. This way, we were able examine the qualitative effects of the controller on the vehicle location distribution.

We begin by discussing the model of the road we use. Next, we describe the control gains used in modeling the vehicle. Subsequently, we detail the parameters used to simulate the measurement dynamics. Finally, we describe the overall structure of the Monte Carlo simulation.

Road model

To examine the behavior of the vehicle on a straight road and on a curve, we ran the simulation on a track that started with a 10m-long segment of straight road. Next, we defined the road to curve through a 180-degree arc with a 10m radius of curvature. The radius of curvature was selected in the range representative of many local roads, as described by [12]–[14]. Following the curved-road section, our simulated road ends with a 10m section of straight road. This way, we can examine both the effects of steady turns and the dynamics of leaving and entering the turns. The map is shown in Fig. 5 where the curved section is in red and the straight section in blue.

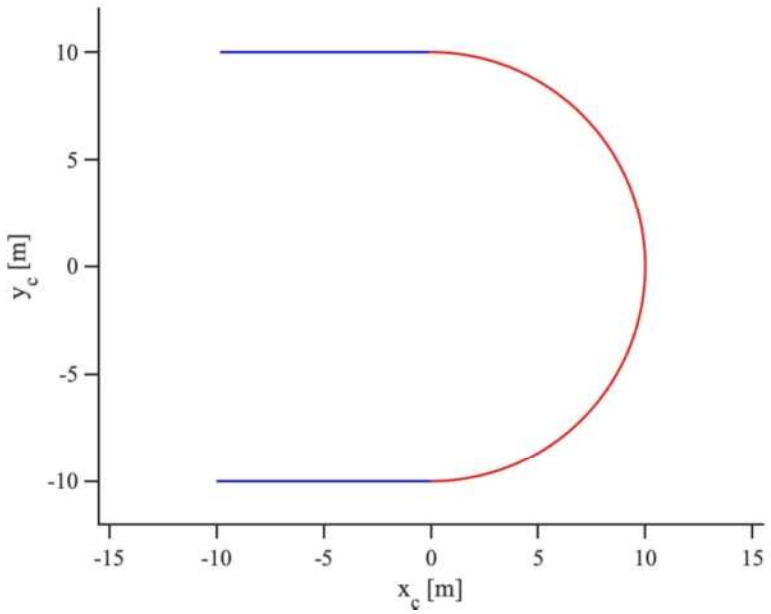

Fig. 5 Nominal Vehicle Trajectory on Curved Road

Car Model

Controller gains for the system can be calculated based on the linear model on a straight road where $\dot{\theta}_{\text{ref}} = 0$. If the angular error $\theta - \theta_{\text{ref}}$ is small, a linear approximation of the system can be made such that (3) becomes

$$\begin{aligned} \dot{x}_p &= v - v_{\text{ref}} \\ \dot{y}_p &= v(\theta - \theta_{\text{ref}}) \end{aligned} \tag{8}$$

By plugging (6) into (8),

$$\begin{aligned} \dot{x} + k_x \mathrm{x} &= -k_x \varepsilon_x + \varepsilon_v \\ \ddot{y}_p + k_\theta v_{\text{ref}} \dot{y} + k_y v_{\text{ref}}^2 y &= -v_{\text{ref}}^2 \left(k_y \varepsilon_y + k_\theta \varepsilon_\theta\right) + v_{\text{ref}} \varepsilon_\omega \end{aligned} \tag{9}$$

where, for the cross-lane direction, $v \approx v_{\text{ref}}$ because the vehicle will be kept close to the target velocity.

For the design of our controller, we intend to have a linear response close to the critically damped response so enable a relatively fast response while also reducing overshoot. From, (9) the damping ratio, ζ, is given by

$$\zeta = \frac{k_\theta}{2\sqrt{k_y}} \tag{10}$$

which serves as a guideline for gain tuning.

Our objective was to create reasonable responses from the vehicle control systems. We expect that a vehicle should respond to roughly within a second. Response times in this range are consistent with models previously used in representing autonomous vehicles[15], [16]. Therefore, in our simulations, we used the gains as shown in Table 1 and the vehicle response is shown in Fig. 6.

Table 1. Controller Gains

k_x	k_y	k_θ
3	$\sqrt{2}$	0.5

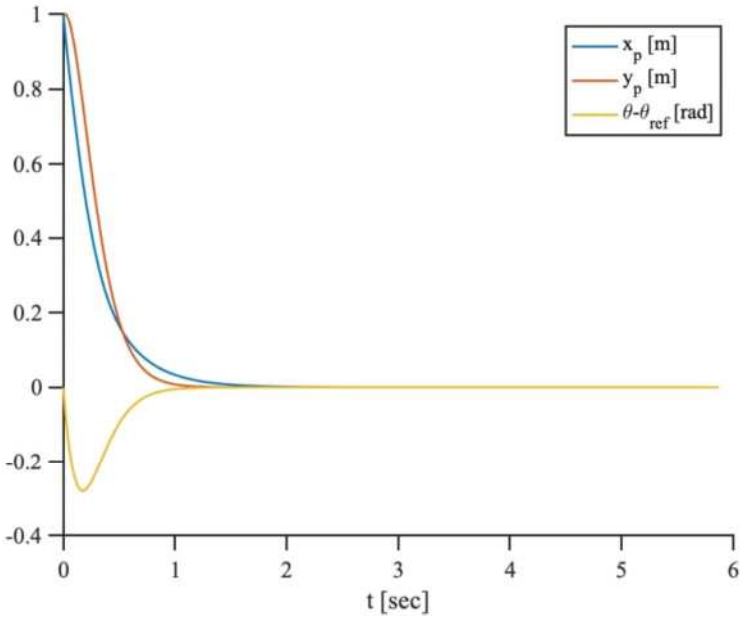

Fig. 6.Vehicle True Errors from an Impulse Response on a Straight Road

Sensor Model

Sensor modeling parameters must also be selected to configure the simulation. Our goal was to tune our sensor noise levels to represent representative GPS/INS fusion systems augmented in the lateral direction by additional perception sensing (e.g. vision or lidar). Since the GPS was assumed to dominate in the along-track direction, we set our one-sigma velocity accuracy to 0.1 m/s and along-track position accuracy to 1.5 m. To account for the perception-aiding in the lateral direction, we set our one-sigma cross-track position accuracy to 0.2 m. To represent a moderate quality INS system, we set the nominal one-sigma heading accuracy to 2 degrees and the nominal angular velocity accuracy to 0.1 deg/s. These standard deviations values are shown in the final row of Table 2.

An important implementation detail involves the measurement correlation model of (7). To explore the impacts of time correlation, we chose $\tau = 10s$ in our GMP model [11]. This time constant represents a challenging but not unrealistic level of GPS multipath error correlation in urban driving. The standard deviations for the input noise to the GPS were tuned (as described in the first row of Table 2) to give the desired output values modeling the sensor noise.

Table 2. Standard Deviation of Sensor Errors

Standard Deviation	σ_x	σ_v	σ_y	σ_θ	σ_ω
GMP Input	2.25m	0.128m/s	0.3m	3deg	0.15deg/s
GMP Output	1.5m	0.1m/s	0.2m	2deg	0.1deg/s

Monte Carlo Simulation

Our simulation is inherently stochastic in that the sensor noise signals are random signals. To capture statistics describing the vehicle position error resulting from our modeled sensor-noise inputs, we implemented a Monte Carlo (MC) simulation with 1000 trials, and computed our results statistically over those 1000 trials. We set a reference 2D route (t, x_{ref}, y_{ref}) defined by initial conditions as well as the reference trajectory shown in Fig. 5, with $v_{\text{ref}} = 10\text{m/s}$.

For each MC trial, we loop through a series of time steps (duration 0.05 s) to a final time $t_N = 5.12\ s$. At each time step of the trial we:

1. Use Table 2 to generate the uncorrelated input noise signals for each first-order GMP described by equation (7), thereby generating the instantaneous correlated error values ε_x, ε_v, ε_θ, ε_y, and ε_ω.
2. Calculate the control outputs at time t_{k} with (6) used to express x_p and y_p. Use gains from Table 1.
3. Use (3) to calculate the state derivatives and propagate states to the next time step, using first-order Euler integration. In other words, find $x_p(t_{k+1})$, $y_p(t_{k+1})$, $\theta(t_{k+1})$, $-\theta_{\text{ref}}(t_{k+1})$).

RESULTS

In this section, we present the results of our investigation. One salient outcome is that the vehicle position distribution is rotated relative to the sensor noise distribution (and to the road). Another notable outcome is that the vehicle's dynamics significantly attenuate the measurement error associated with the heading sensor.

To visualize the first of these phenomena, associated with the vehicle's position distribution, consider Fig 7. The figure provides a conceptual illustration of the vehicle position distribution at two locations: along the straight section of the road and along the curve. As illustrated, the elliptical distribution of vehicle positions (cyan) aligns with the roadway (yellow) on the straight section, but not on the curved section. Considering that the vehicle is moving along the path from left to right, the angle of the rotation of the error ellipse lags the reference heading angle.

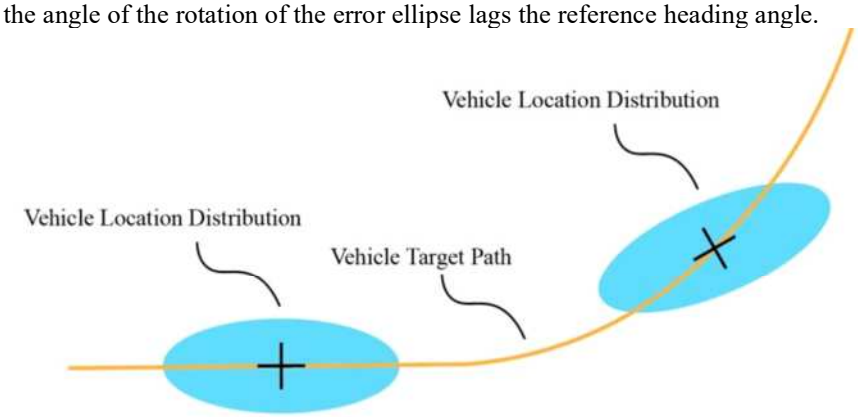

Fig. 7. Orientation of Vehicle Location Probability Distribution on Curved and Straight Roads

The conceptual illustration is useful to interpret the statistical data, which are shown in Fig. 8. The MC-generated position errors on the straight section (blue dots, top of Fig. 8) are distributed in a manner that aligns along the roadway.

The MC-generated position errors on the curved section (red dots, bottom of Fig. 8) are distributed in a manner that does not align with the roadway. Note that the curved roadway on the bottom of Fig. 8 has been rotated so that the local tangent direction is horizontal at the reference point, which is centered at the coordinate (0,0). Thus, the fact that the red position-error distribution is not aligned in a horizontal direction indicates that it is not aligned with the instantaneous reference heading.

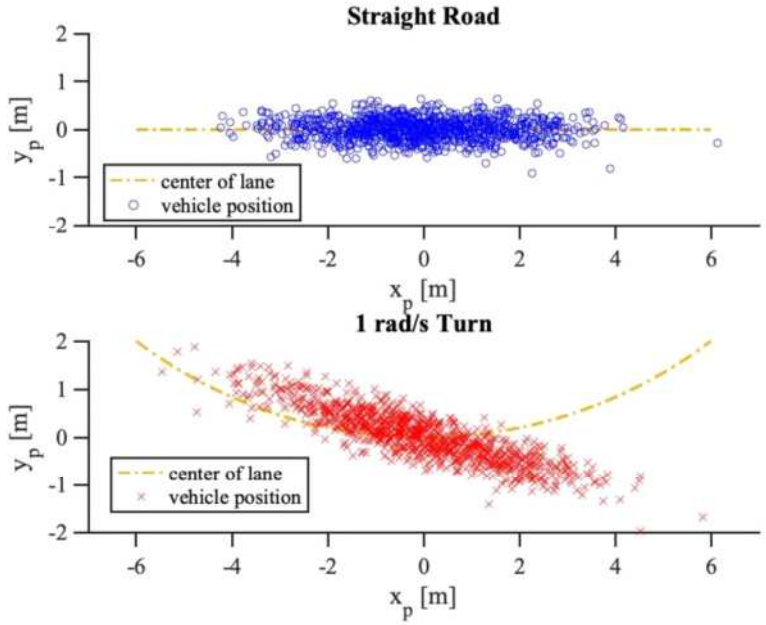

Fig. 8. Position Errors on a Curved and Straight Road

The second notable phenomenon involves the vehicle heading distribution. In our simulation, heading-sensor measurement errors are significantly attenuated by the control system and vehicle dynamics. For instance, this can be observed by comparing the error-distributions standard deviations for a specific time step in the MC simulation. Standard deviations for two specific time steps (at steady-state on the straight and at steady-state on the turn) are compiled in Table 3. During steady maneuvers, the vehicle heading errors, as compared to the reference $\theta_{\text{ref}}(t)$, are generally much smaller (with one-sigma values of 0.45° on a straight turn and 0.97° on the turn) than the heading-sensor noise (with a one-sigma value of 2°). The exception is on the transitions between the straight and curved sections, where the controller transient error has not converged. Our simulation resulted in a one-sigma vehicle heading error of 6.45° during this transition, which was significantly larger than the measurement error of 2.0°.

Table 3. Outputs from Monte Carlo Simulation

Standard Deviation	σ_x	σ_v	σ_y	σ_θ	$\sigma_{\theta loc}$	σ_ω
GMP Output	1.50 m	0.1 m/s	0.20 m	2.0 deg	N/A	0.1 deg/s
Vehicle State Straight	1.70 m	N/A	0.24 m	0.45 deg	0.45 deg*	N/A
Vehicle State Turn	1.68 m	N/A	0.54 m	9.71 deg	0.97deg*	N/A

** Taken relative to local road heading*

For safety analysis, it is relevant to characterize heading error relative to the local lane boundaries. Consider the case when the vehicle is lagging the current reference point on the nominal trajectory. During the turn, the direction of the lane boundaries at the reference point (which is ahead of the vehicle) will be different than the direction of the lane boundaries nearest the vehicle. For this reason, we introduce the variable θ_{loc} which describes the orientation of the

road and lane boundaries, not at the reference location, but at the place along the road closest to the vehicles true position. The one-sigma heading errors relative to θ_{loc} are significantly smaller than those relative to the reference points. In fact, the one-sigma heading errors are even smaller than the navigation system true errors. In our simulation, the 2.0 degrees standard deviation of the sensor model was reduced to 0.97 degrees as shown in column 5 of Table 3.

The difference in the heading accuracy as defined relative to θ_{loc}, as opposed to the reference point heading θ_{ref}, can be interpreted in terms of Fig. 9, which shows heading error values on the straight road (blue dots) and curved road (red dots), plotted as a function of the longitudinal error. Note that θ_{loc} is also plotted (yellow dash-dot line). On the curved road, the figure shows that the local heading error relative to θ_{ref} is highly correlated with θ_{loc}, which varies linearly along the longitudinal coordinate.

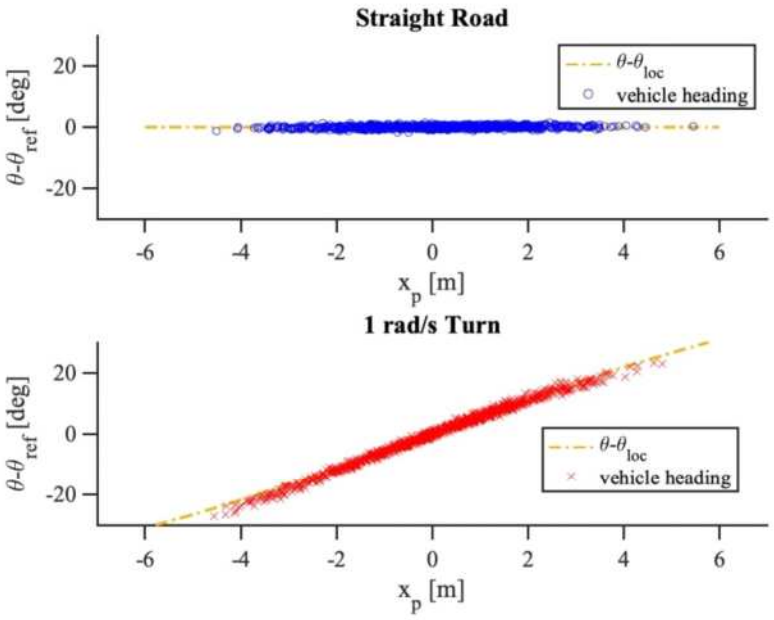

Fig. 9. Longitudinal and Heading Errors on a Straight Road (top) and Curved Road (bottom)

The situation is quite different in the transition region from the straight to the curve. To understand this region, where turn initiation is dominated by the open loop controller, it is instructive to consider the differences between the heading error relative to the nominal trajectory and the heading error relative to the road. This can be seen in Fig. 10. The figure shows the transition region, but it is otherwise analogous to Fig. 9, plotting heading errors as a function of longitudinal errors. The yellow reference line shows the sudden transition in reference heading from the straight road, where the reference is horizontal, to the curved section, where the reference has a positive slope. When the longitudinal error is positive, the vehicle has entered the curve early and the heading error is large.

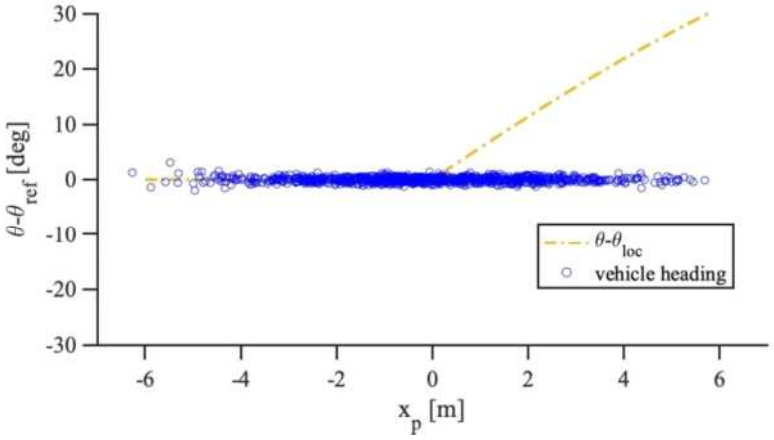

Fig. 10. Lateral Error and Local Orientation Error at the Beginning of a Turn

It is also instructive to consider the full time-history of the statistics generated by the MC simulation, and not just the three points considered so far (i.e., steady-state on the straight section, steady-state on the curve, and the transition point between the straight and curved sections). To create a continuous picture of the errors as a function of time, we focus on the cross-correlation coefficient between scalar errors and specifically ρ_{xy} (the correlation between x_p and y_p) as well as $\rho_{x\theta}$ (the correlation between x_p and $\theta - \theta_{\text{ref}}$). The correlation coefficients provide a compact summary of the previous figures. For instance, in Fig. 9 it is clear that $\rho_{x\theta} \sim 0$ for the straight section (where the longitudinal and heading errors are uncorrelated), but that $\rho_{x\theta} \sim 1$ for the curved section (where the longitudinal and heading errors are entirely correlated). By comparison, Fig. 10 lies somewhere in the middle, with $0 < \rho_{x\theta} < 1$. To analyze the full time-history, we plot correlation coefficients as a function of time, computed over all of the MC trials (see Fig. 11). In this figure, the domain where the reference is on a curved path is plotted in red, and the domain on the straight path is plotted in blue.

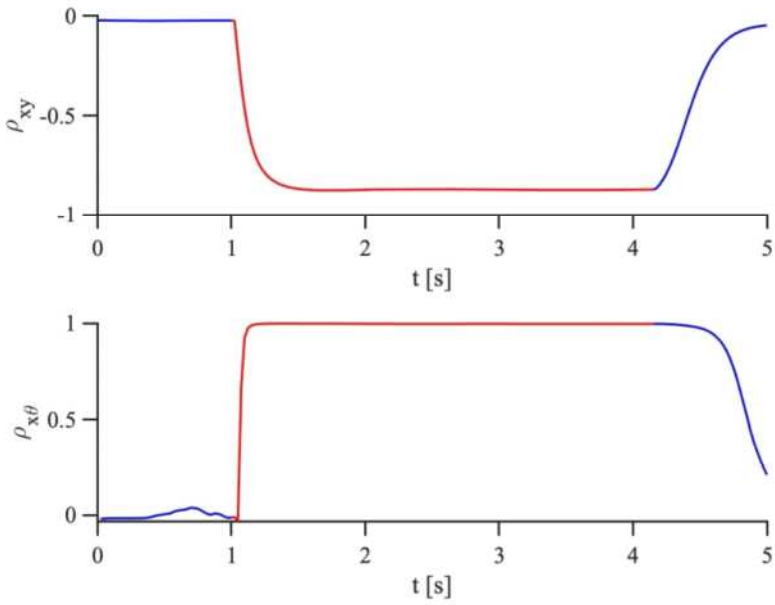

Fig. 11 Error Covariances of Each Timestep of the Simulation on the Curved Road

As a final analysis, we also plot the correlation coefficient $\rho_{x\theta_{loc}}$, which describes the correlation between the longitudinal error and the local heading error $\theta - \theta_{loc}$ (see Fig. 12). The vehicle correlation coefficient in the curve is

much smaller for the local-heading error of Fig. 12 as compared to the reference-heading error of Fig. 10 (bottom). Observe that the local-heading errors are greatest at the begging and end of the turns.

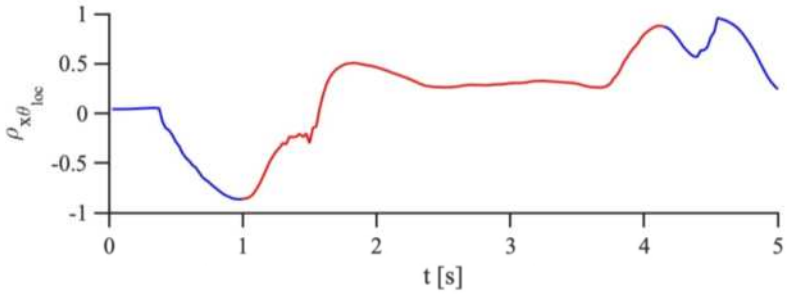

Fig. 12. Error Covariance of True Lateral Error and Local Orientation Error

DISCUSSION

The results presented in the prior section qualitatively describe the stochastic dynamics of the system on straight and curved roads. The two primary results were that (i) the vehicle position-error ellipse is not aligned with the road direction, at least on curved roads, and that (ii) the vehicle dynamics generally mitigate heading errors. Here, we will explain the causes of these results, contextualize them, and suggest their significance.

Position Error (on Curved Roads)

On curved roads, the, along-path and cross-path error distributions become correlated. This correlation can lead to a significant increase in the lateral position error of the vehicle relative to the road center as shown in Fig. 8. In this case, the lateral errors are nearly doubled by the rotation of the error distribution, and the chance of a lane excursion would likely be significantly increased. This likelihood of a lane excursion is particularly prevalent if the vehicle's true position is ahead of its observed position as the road curves away, as can be seen by the road direction in Fig. 8.

These effects should be considered in safety analyses, for instance in deriving protection levels and comparing them to alert limits, as discussed in [1], [2]. Both papers assumed protection-limits are generated directly from sensor error distributions and compared to alert limits generated from the geometry of the road. This approach misses a key detail, in that the major and minor axes of the vehicle position distributions do not necessarily align with the road or the sensor error distribution. Thus, our findings in this paper suggest that the existing methods of [1], [2] might need to be adapted to address, for instance, the rotated distributions shown in Fig. 8.

It is interesting to think about why the position uncertainty distribution becomes rotated in a curve. From a mathematical point of view, the rotation is driven by the cross product in equation (3), where rotation is related to the size and direction of the error.

With regards to safety analyses, we also observe a second, relatively minor effect, in that the distribution of vehicle-position errors is slightly larger than the distribution of position-sensor errors, as shown in Table 3. We can attribute this discrepancy to the fact that we used a proportional controller in our model. Therefore, we expect that persistent (time-correlated) sensor errors in vehicle speed and orientation would generate steady-state true errors in position. It may be possible to mitigate these effects with control tuning (such by the addition of an integrator), but it is important to be cautious about over-turning the control model in this analysis, since optimizing in one context may degrade another.

Heading Error (on Straights and Curved Roads)

Our investigation observed that vehicle-heading errors were significantly smaller than sensor-heading errors, at least on steady straight and curved roads. We believe sensor-heading errors are reduced because they quickly accumulate as lateral error y_p, which is detected by lateral sensors and countered by the controller. Therefore, any amount of

heading error quickly becomes observable to the controller such that the heading error is quickly mitigated by the system. This explanation accounts for the 5-fold decrease in heading errors shown in Table 3 on a straight or curved road.

For safety applications, the most important expression of heading error is the error relative to closest point on the reference trajectory (the local road) and not relative to the current reference point, since the vehicle may race ahead or fall behind the moving reference point. In defining an alert limit, for example, the size of the alert limit is influenced by the orientation of the vehicle relative to the local road section [1], [2]; vehicle heading errors tend to erode safety margins around a vehicle. Fortunately, the orientation errors relative to the local road are quite small in the steady straight and curved road cases, which greatly reduces the penalty associated with uncertain heading.

An interesting observation is that the small local-heading errors can be explained as a necessary sign that the controller is working effectively. An effective controller will keep the vehicle from deviating too far away from the path. From (3), we observe that a necessary step to achieve this is to keep the vehicle oriented in the direction of the path. Therefore, this observation is the result of the controller, and shows that the heading error is largely related to the road curvature and how far the vehicle is along the road.

In our simulations, we see that the vehicle tends to miss the timing of the transition if there is an unobservable along path position error. This effect is closely related to the guidance, which we modeled using a feedforward term, which are key to triggering the vehicle rotation on the curved road section. If the vehicle un-observably drifts slightly ahead of the reference point, then the guidance system triggers too late and the vehicle continues straight past the transition point, as shown in Fig. 10. The shape of the distribution in Fig. 10 represents the case at the moment just before the transition, when the reference angular velocity has not yet changed; the distribution shape is very different shortly after, when the reference guidance becomes active. This behavior is expected to be less extreme on a real vehicle, as features at a corner might become observable and augment the navigation errors, so it would be necessary to refine the navigation and controller model to obtain a representative analysis of such transitions.

Future Work

The obvious and necessary next step is to build this analysis into defining protection levels and alert limits. Whereas references [1], [2] directly relate sensor errors to road geometry, this paper clearly shows that the vehicle dynamics transform sensor errors, and so sensor errors should not be assumed to be identical to vehicle state errors when developing protection levels. Developing a thorough approach to bound the effects of vehicle dynamics on sensor errors will be challenging, requiring a generalization of the approach used in this paper, where we consider only one very specific scenario. One step is to consider a range of vehicle types and roadways; another is to consider a range of possible control laws and control parameters; a third is to consider a range of possible sensor-noise models, including a range of correlation times. Another important consideration is to include a look-ahead period, since the protection level analysis needs to be valid not just at the current time, but also for some period into the future. Such an analysis will require new analytical technique to project when combined errors (e.g., position and heading) create the largest risk of a lane excursion.

To support future safety analyses, another important step will be to validate the use of abstracted models like the ones presented in this paper. Though we believe that the simplified model is generally representative of a wide range of real-world controllers, it is not clear how accurately our approach represents autonomous vehicles being deployed and tested in the field. Automated lane-keeping experiments could help address these questions, potentially for operations involving a range of errors sources, terrains, and environmental conditions.

SUMMARY

In this paper, we considered how vehicle system dynamics map measurement errors into vehicle state errors. To this end, we conducted a Monte Carlo simulation of an autonomous vehicle following a path with random time-correlated sensor errors in position and orientation. We showed that errors in the sensing of vehicle heading do not cause large errors in true vehicle heading, since vehicle-heading errors become observable indirectly, manifesting as cross-track position errors. Furthermore, we showed that, when a vehicle rounds a curve, its position error distribution may not

align with the road, since time-correlation causes apparent lag. These results suggest that measurement errors should not be compared directly to alert limits generated from road and vehicle geometry, but rather that the effects of system dynamics on the measurement error should also be considered.

ACKNOWLEDGMENTS

The authors wish to acknowledge and thank the U.S. Department of Transportation Joint Program Office (ITS JPO) and the Office of the Assistant Secretary for Research and Technology (OST-R) for sponsorship of this work. We also gratefully acknowledge NSF grant CNS-1836942, which supported specific aspects of this research. Opinions discussed here are those of the authors and do not necessarily represent those of the DOT, NSF, or other affiliated agencies.

REFERENCES

[1] T. G. R. Reid *et al.*, "Localization Requirements for Autonomous Vehicles," *SAE Int. J. Connect. Autom. Veh.*, vol. 2, no. 3, pp. 1–16, 2019.

[2] O. N. Kigotho and J. Rife, "Comparison of Rectangular and Elliptical Alert Limits for Lane-Keeping Applications," in *ION GNSS+*, 2021.

[3] B. DeCleene, "Defining Pseudorange Integrity - Overbounding." pp. 1916–1924, 22-Sep-2000.

[4] J. Rife and S. Pullen, "The Impact of Measurement Biases on Availability for CAT III LAAS," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 42, pp. 1386–1395, Jun. 2005.

[5] J. Rife, S. Pullen, P. Enge, and B. Pervan, "Paired overbounding for nonideal LAAS and WAAS error distributions," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 42, no. 4, pp. 1386–1395, 2006.

[6] J. D. Larson, D. Gebre-Egziabher, and J. H. Rife, "Gaussian-Pareto overbounding of DGNSS pseudoranges from CORS," *Navig. J. Inst. Navig.*, vol. 66, no. 1, pp. 139–150, Mar. 2019.

[7] J. H. Rife, "Robust Chi-square Monitor Performance with Noise Covariance of Unknown Aspect Ratio," *Navig. J. Inst. Navig.*, vol. 64, no. 3, pp. 377–389, Sep. 2017.

[8] P. D. Groves, Z. Jiang, M. Rudi, and P. Strode, "A portfolio approach to NLOS and multipath mitigation in dense urban areas," *26th Int. Tech. Meet. Satell. Div. Inst. Navig. ION GNSS 2013*, vol. 4, no. September, pp. 3231–3247, 2013.

[9] P. Xie and M. G. Petovello, "Measuring GNSS Multipath Distributions in Urban Canyon Environments," *IEEE Trans. Instrum. Meas.*, vol. 64, no. 2, pp. 366–377, 2015.

[10] N. S. Nise, *Control Systems Engineering*, 7th ed. Wiley, 2015.

[11] M. G. Petovello, K. O'Keefe, G. Lachapelle, and M. E. Cannon, "Consideration of time-correlated errors in a Kalman filter applicable to GNSS," *J. Geod.*, vol. 83, no. 1, pp. 51–56, 2009.

[12] *A Policy on Geometric Design of Highways and Streets*. American Association of State Highway and Transportation Officials, 2001.

[13] *Roundabouts: An Informational Guide*, vol. FHWA-RD-00-067. US Department of Transportation Federal Highway Administration.

[14] *Design Manual*. Washington State Department of Transportation, Olympia, WA., 2011.

[15] J. Jiang and A. Astolfi, "Lateral Control of an Autonomous Vehicle," *IEEE Trans. Intell. Veh.*, vol. 3, no. 2, pp. 228–237, Jun. 2018.

[16] Ş. Y. Gelbal, S. Tamilarasan, M. R. Cantaş, L. Güvenç, and B. Aksun-Güvenç, "A connected and autonomous vehicle hardware-in-the-loop simulator for developing automated driving algorithms," *2017 IEEE Int. Conf. Syst. Man, Cybern. SMC 2017*, vol. 2017-January, pp. 3397–3402, Nov. 2017.