Sensor Fusion and Bang-Bang Control with Nonholonomic Constraints*

Douglas LYON**

A bang-bang control law is used to provide guidance for a car's steering system. The control law uses a proportional plus derivative error function which operates on the output of an observer and is suitable for embedded control of a car. The observer fuses sensor data from sensors whose reliability is inversely related to their availability. An erroneous plant model is present and is subject to nonholonomic and kinematic constraints. The sensors measure the plant with 25% systematic error at 100 hertz. Internal sensors are sampled at 20 hertz and have 10% error. External sensors are sampled at 5 hertz and have no error. The observer performs dead reckoning during the intervals of time when the external sensors are not available. There are two internal sensors, an odometer and a steering-wheel angle meter. The odometer is used to compute a linear approximation to the speed of the plant. The steering-wheel angle meter measures the steering-wheel deflection relative to the longitudinal centerline of the plant. There are three external sensors which are used to measure the pose of the plant.

Key Words: Nonholonomic Systems, Bang-Bang Control, Sensor Fusion, Embedded Control

1. Introduction

The problem is to automate sensor fusion and to control a car such that it will track a reference path. We assume that the car has sensors, effectors and sufficient on-board computing power to run a tracking program in real time.

We assume that systematic error is a worst-case error. We also assume that there are two kinds of sensors, internal and external, and that these sensors have an availability which is inversely related to their reliability. The sensor's sample rates are known.

We assume that the car steers with its front wheel and that it has known limits on turning radius and steering acceleration. It is assumed that the front of the car follows the front wheel without slip, and so the car is subject to nonholonomic constraints. We assume that the rear wheel remains parallel with the main axis of the car. The objective is to get the rear of the car to track a reference path. We do not model the differential coupling in the drive train or the width of the car.

The problem is: given a car with discrete sensor feedback, and modeling error in steering inertia, subject to nonholonomic and kinematic constraints (curvature and acceleration), develop a program that is able to track a reference path.

The problem of sensor fusion using nonholonomic constraints has significant implications for factory-floor motion planning. For example, the use of automated vehicles for moving parts is becoming common. These vehicles follow predefined paths that are inflexible. In addition autonomous-motion planning is a requirement of vehicles which cannot be operated remotely, as happens over great distances (i.e., a planetary rover). Reference path tracking with nonholonomic constraints can also be a product for automating the parallel parking of a car[1].

2. Open-Loop Control

Open-loop control establishes upper and lower bounds on controller performance by modeling an observer with and without plant-model error. Figure

* Received 15th January, 2001

** Chair of Computer Engineering Department, Fairfield University, Fairfield, CT, 06430, U.S.A. E-mail: lyon @docjava.com

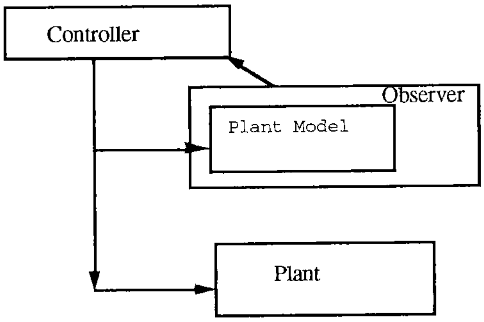

Fig. 1 Open-loop control. The observer uses controller-input to estimate plant states. The controller uses the plant-state estimate to formulate input.

1 shows an overview of an open-loop system with three main parts: the controller, observer and plant. We assume that the plant model has error.

The controller uses the plant-state estimation produced by the observer to determine the control signals. These signals are applied to the plant model and the plant. Error in the plant model will cause a reaction to the control signals which is different from the plant's reaction. This is open-loop control because there are no sensors between the plant and the observer. The observer will perform sensor fusion to update the plant model in the closed-loop control regime of Section 3. The difference between closed-loop and open-loop control is the presence of sensor feedback between the plant and the observer, otherwise the control regimes are identical.

2.1 Nonholonomic constraints and the reference path

The open-loop controller models the steering system as having rotational inertia and alters steering acceleration to control steering-wheel angle. Let (x, y) be the position of the front of the plant model, (x_r, y_r) be the position of the rear of the plant model, s be the length of the path traveled by the plant model, θ be the plant-model orientation, angle formed by the main axis of the plant with the X-axis of a Cartesian coordinate frame aligned with the curb, ϕ be the steering-wheel angle relative to the longitudinal centerline of the plant model and L be the length of the plant and the plant model.

The temporal differential equations which make up the nonholonomic constraints are

$$\frac{dx}{dt}=\frac{ds}{dt}\cos(\theta+\phi) \quad (1)$$

$$\frac{dy}{dt}=\frac{ds}{dt}\sin(\theta+\phi) \quad (2)$$

$$\frac{d\theta}{dt}=\frac{ds}{dt}\frac{\sin\phi}{L} \quad (3)$$

For a proof see (2). The car pose (x, y, θ) and

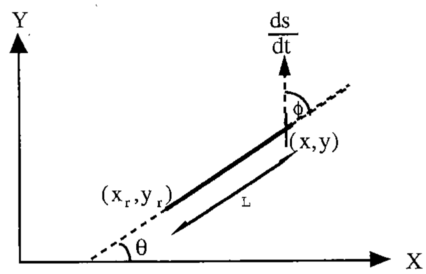

Fig. 2 Car notation. The front of the car follows the front wheel and this causes the nonholonomic constraints.

steering-wheel angle (ϕ) are shown in Fig. 2.

While the reference path may take on any form, in this particular case we shall use a curvature constrained fifth-order polynomial

$$y5(x_r)=y_e\left[6\left(\frac{x_r}{x_e}\right)^5-15\left(\frac{x_r}{x_e}\right)^4+10\left(\frac{x_r}{x_e}\right)^3\right] \quad (4)$$

whose curvature is

$$k5(x_r)=y_e\left(120\frac{x_r^3}{x_e^5}-180\frac{x_r^2}{x_e^4}+60\frac{x_r}{x_e^3}\right)$$
$$\times\left[1+y_e^2\left(30\frac{x_r^4}{x_e^5}-60\frac{x_r^3}{x_e^4}+30\frac{x_r^2}{x_e^3}\right)^2\right]^{-3/2} \quad (5)$$

For a proof see (3).

2.2 Bang-bang control with nonholonomic constraints

This section examines the plant model and controller. The goal is to formulate a control law which enables the car to track the curvature of the reference path. The controller specifies angular acceleration to the steering system. The reference path's curvature is a function of x_r. This, in turn, is a function of x and θ. These state-variables are subject to the nonholonomic constraints of (1), (2) and (3). Let the reference path's curvature be denoted by $k5$. Curvature is a function of the rear-component of the X-coordinate of the car. The relationship between the front and rear coordinates is given in

$$x_r=x-L\cos\theta$$
$$y_r=y-L\sin\theta \quad (6)$$

The bang-bang control law is given by

$$a_\phi=\begin{cases} a_{\phi\,\max} & \text{for } e<0 \\ -a_{\phi\,\max} & \text{otherwise} \end{cases} \quad (7)$$

where

$$a_\phi=50\ \text{rad/s}^2 \quad (8)$$

The decision variable, e, is an error signal computed by

$$e=\phi-\phi5+\alpha\left(\frac{d\phi}{dt}-\frac{d\phi5}{dt}\right)$$
$$+\alpha_\theta\left[\theta-\theta5+\alpha\left(\frac{d\theta}{dt}-\frac{d\theta5}{dt}\right)\right] \quad (9)$$

where

The mix ratio between the proportional and derivative control must be found by experiment. The values

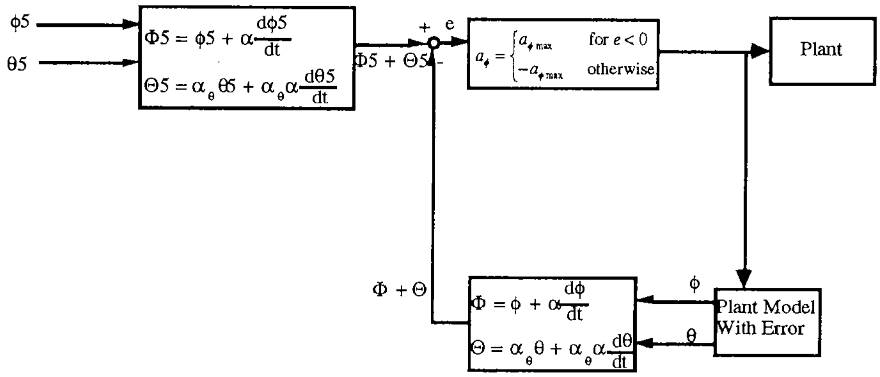

Fig. 3 Proportional plus derivative controller

$$\alpha_\theta=2 \text{ and } \alpha=0.05 \quad (10)$$

gave good results.

The reference steering-wheel angle of the car is

$$\phi5=\arctan(k5L) \quad (11)$$

and reference orientation is

$$\frac{d\theta5}{dt}=\frac{ds}{dt}\frac{\sin(\phi5)}{L} \quad (12)$$

3. Closed-Loop Control

This section summarizes an investigation into closed-loop control, and this is compared with the open-loop control of Section 2. Closed-loop control improves performance by using sensor feedback to improve the observer's plant-state estimates.

Section 3.1 provides an overview of the closed-loop system. Section 3.2 covers internal and external sensors with a justification for sensors whose reliability is inversely related to availability. Section 3.3 summarizes closed-loop bang-bang control for 2-D motion using internal and external sensors.

3.1 System overview

Sensor feedback is used by the observer to compensate for the error in the plant model. Figure 4 shows an overview of a closed-loop system with three main parts: a controller, an observer and a plant.

The observer uses a plant model and sensor measurements to estimate the state of the plant. The process that the observer uses to incorporate measurements from different sensors is called *sensor fusion*. There are two classes of sensors to be fused, internal and external. Their reliability is inversely related to their availability. External sensors are more accurate than internal sensors and dominate the internal sensors when they are available.

Internal sensors measure the plant state relative to itself. An odometer is an example of an internal sensor which measures the distance a car travels by measuring tire motion. Since tires slip, the odometer will have some error. We also assume that internal sensors have a 20 hertz sampling rate. Candidate internal sensors are: the accelerometer, the tachometer and the odometer. These sensors give equivalent data. Equivalent data means that data provided by any one of the sensors may be transformed into data provided by any of the others.

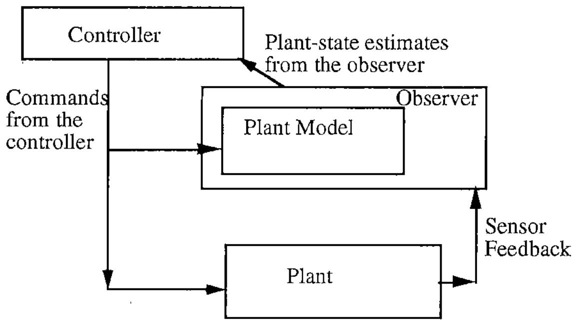

Fig. 4 Closed-loop control. The observer uses the plant model driven by the controller commands to estimate plant states. The controller uses the plant-state estimates to generate commands. Sensor feedback is used to correct the observer's estimate.

External sensors measure the plant state relative to the environment. A Sonar sensor is an example of an external sensor. We assume that external sensors have a 5 hertz sampling rate.

3.2 Sensor fusion

In this section we describe a technique for fusing the information from the internal and external sensors using the nonholonomic constraints on orientation. A block diagram is shown in Fig. 4.

This section summarizes a study of 2-D control using the fifth-order polynomial to generate a reference path. Bang-bang control is used to adjust the acceleration and the steering-wheel angle of the car. Nonholonomic constraints are used by the observer to compute pose (position and orientation).

We compare open-loop control against several closed-loop control regimes. These regimes differ in the sensor data provided to the observer. Open-loop control is based on a plant-model with error and does not use sensors. The closed-loop control regimes are based on internal sensors (dead reckoning), external sensors and a combination of both. As a basis for comparison we also show control based on error-free sensing of the plant states. This provides a best-case limit on the parking performance and can be compared with open-loop control (which provides a worst-case limit on parking performance).

External sensors are sampled without error at 5 hertz and internal sensors are sampled with 10% error at 20 hertz. The plant-model is used between sensor readings and is always available to the controller. The plant-model has 25% steering and linear acceleration error. Sensors compensate for the error in the plant-model, improving performance.

We add sensors for steering-wheel angle and car-orientation to enable 2-D dead-reckoning. It is assumed that we are able to sense the steering-wheel angle (ϕ_p) and the odometer (s_p) with 10% error and with 20 hertz sampling. The sensors consistently under-report the true magnitude of the sensed variables.

The reference steering-wheel angle is a function of the fifth-order polynomial's curvature (as shown in Section 2). The controller is able to change the steering-wheel angle by banging on the steering system. The control law described in Section 2 uses values measured by the internal sensors. These are multiplied by 0.90 to model the 10% measurement errors of the internal sensors.

We use ultrasonic ranging for pose (position and orientation). We assume that the external sensors are able to sample at a rate of five hertz and that they have no error. During the intervals in which external sensors are not available, the plant-model is updated by the internal sensors. It will be shown that fusion with external sensors is not much better than fusion with internal sensors and that fusion with both internal and external sensors is better than fusion with either one alone.

The plant model has a 25% error in linear and steering acceleration. The observer fuses the internal and external sensors to reduce the plant-model error. The plant steering-wheel angle is measured directly at 20 hertz with 10% error. This is used to improve the observer's estimate of the steering-wheel angle. External sensors cannot help improve the observer's estimation of the steering-wheel angle because the change in the plant's pose lags behind the change in the steering-wheel angle. Car orientation is directly measured by an external sensor, without error, and is used to update the observer's estimate.

During the intervals of time when the externally sensed plant orientation is not available, the observer

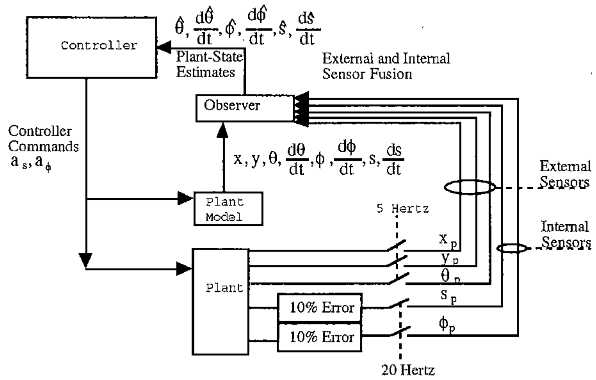

Fig. 5 Internal and external sensors used by the observer. Internal sensors are sampled at 20 hertz and have 10% error while external sensors are sampled at 5 hertz and have no error.

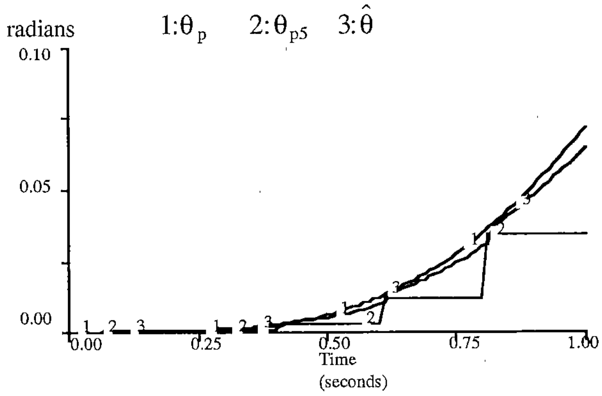

Fig. 6 The Observer's use of the external orientation sensor. The plant's orientation (θ_p) is sampled accurately at 5 hertz by the orientation sensor whose output is denoted θ_{p5}. In between sensor samples the observer uses (13) to compute an estimate of the plant's orientation ($\hat{\theta}$).

uses the nonholonomic constraints on orientation to incorporate the internally sensed estimate of the steering-wheel angle with the fused estimate of speed. Thus

$$\frac{d\hat{\theta}}{dt}=\frac{d\hat{s}}{dt}\frac{\sin(\hat{\phi})}{L} \tag{13}$$

is the rate of change of orientation computed by fusing information from internal sensors. Orientation is obtain via numeric integration of (13) and exhibits cumulative error which is corrected by the external orientation sensor. This is shown in Fig. 6.

Similarly, reference orientation is computed using

$$\frac{d\theta 5}{dt}=\frac{d\hat{s}}{dt}\frac{\sin(\phi 5)}{L} \tag{14}$$

Finally, the difference signals used to compute when to bang on the steering combine (13) and (14) which results in

$$e_{\frac{d\theta}{dt}}=\frac{d\hat{\theta}}{dt}-\frac{d\theta 5}{dt} \tag{15}$$

This control law is just like that of Section 2, except for the introduction of the sensor feedback for improving the plant-state estimates.

The fused estimate of speed is the result of numerically differentiating the odometer reading at 20 hertz to obtain a tachometer reading with error, adjusted by the derivative of the positional change using the external sensors. The observer's estimate of the speed of the plant diverges from the plant's true speed. This is due to the 10% odometry error which is used to computer the plant's speed during the interval when external sensors are unavailable. The odometry-based speed estimation is independent of the steering wheel angle estimation in the non-holonomic constraints.

3.3 Simulation results

In this section we compare the open-loop control results with the closed-loop control results. Figure 7 shows a comparison of control regimes. Since there is no single objective function to optimize, several plant states and observer estimates are shown. The observer estimates are shown in order to establish that the controller is performing optimally given the observer's information. As the information to the observer improves, so does its plant-state estimation and the resulting controller performance.

In the open-loop control with error in the plant model, we established the worst-case bound on controller performance. Open-loop control is based on a plant-model with error and does not use sensors. This was discussed in Section 4 where engine and steering acceleration were underestimated by 25%. The overshoot of the plant, when the speed is zero, is a useful criterion and is 3.76 m−2.41 m=1.35 m. Better observer estimates reduce this overshoot.

The error-free sensing of the plant, with 25% error in the plant model, sets a lower best-case performance bound for closed-loop control. Sensors are error-free and running at the simulator rate of 100 hertz. Without using an adaptive reference model, and with no a priori knowledge of the plant-model error, this must be the best performance that we can expect from any sensor-fusion scheme. The overshoot using this scheme is 2.77 m−2.41 m=0.36 m.

The closed-loop plant with internal sensors has 25% error in the plant-model and 10% error in the internal sensors. Also, the internal sensors sample at 20 hertz. The overshoot of the plant position when the speed is zero is 3.27 m−2.41 m=0.86 m.

The closed-loop plant with external sensors has 25% error in the plant-model and no error in the external sensors. Also, the external sensors sample at 5 hertz. The overshoot of the plant position when the speed is zero is 3.21 m−2.41 m=0.80 m.

	Position When the Brakes are Applied	Speed When the Brakes are Applied	Time When the Brakes are Applied	Position when speed is zero.	Orientation when speed is zero	Front-wheel angle when speed is zero.	Time When Speed is Zero
Error-Free Plant Model (a)	1.51 m	1.59 m/s	1.91 s	2.41 m	0 rad	0 rad	3.04 s
Error-Free Sensing of the Plant with 25% Error in Plant Model (c)	1.39 m	1.70 m/s	1.64 s	2.77 m	.005 rad	-.004 rad	3.26 s
Open-Loop Plant Model with error (b)	1.51 m	1.59 m/s	1.91 s	2.41 m	0 rad	0 rad	3.04 s
Open-Loop Plant	1.88 m	1.98 m/s	1.91 s	3.76 m	-.46 rad	-.039 rad	3.79 s
Closed-Loop Observer with Internal Sensors	1.47 m	1.63 m/s	1.78 s	2.90 m	.0075 rad	-.0034 rad	3.54 s
Closed-Loop Plant with Internal Sensors (d)	1.63 m	1.85 m/s	1.78 s	3.27 m	.0073 rad	0 rad	3.53 s
Closed-Loop Observer with External Sensors	1.41 m	1.69 m/s	1.77 s	3.10 m	.0225 rad	-.04 rad	3.57 s
Closed-Loop Plant with External Sensors (e)	1.60 m	1.84 m/s	1.77 s	3.21 m	.0227 rad	-.0338 rad	3.51 s
Closed-Loop Observer with Internal and External Sensors	1.40 m	1.68 m/s	1.74 s	2.79 m	-.0315 rad	-.00838 rad	3.60 s
Closed-Loop Plant with Internal and External Sensors (f)	1.56 m	1.81 m/s	1.74 s	3.16 m	-.0314 rad	-.00375 rad	3.47 s

Fig. 7 Comparison of control regimes. The open-loop plant shows the worst-case response with 1.35 m overshoot. The error-free sensing of the plant, with 25% error in the plant model, sets a best-case performance bound for closed-loop control with 0.36 m overshoot. The 2-D closed-loop plant with internal sensors has a 0.86 m overshoot. The closed-loop plant with external sensors has a 0.80 m overshoot. The closed-loop plant with internal and external sensors has a 0.75 m overshoot.

The closed-loop plant with internal and external sensors has 25% error in the plant-model. There is no error in the external sensors, which sample at 5 hertz. There is 10% error in the internal sensors which sample at 20 hertz. The overshoot of the plant position when the speed is zero is 3.16 m−2.41 m=0.75 m.

In the closed-loop plant with internal and external sensors, the performance after a single maneuver is slightly better than using internal sensors alone. An examination of the plant's position when the speed is zero shows that the fusion of external sensors with internal sensors will reduce overshoot by 0.11 meters, for a single maneuver. It is also seen that internal sensors reduce the single-maneuver overshoot obtained from open-loop control by 0.49 meter.

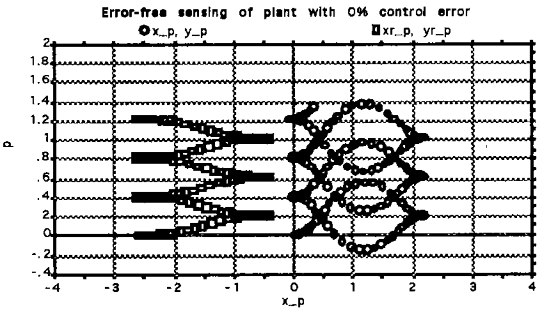

(a) Error-free sensing of plant with error-free plant model

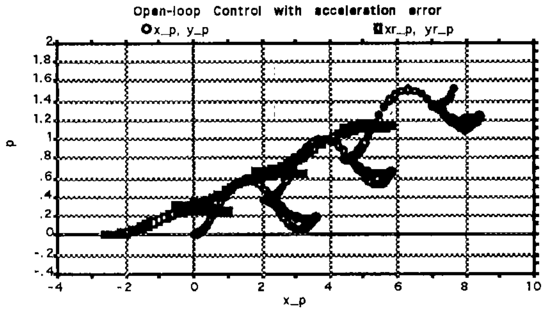

(b) Open loop

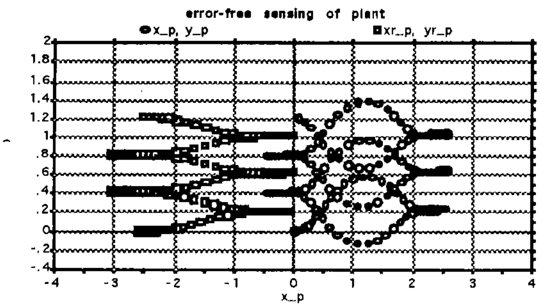

(c) Error-free sensing of the plant

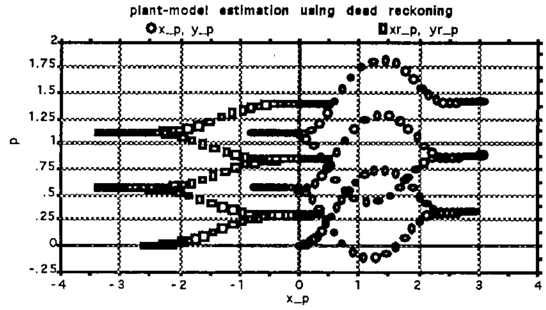

(d) Internal sensors with 25% error in plant model

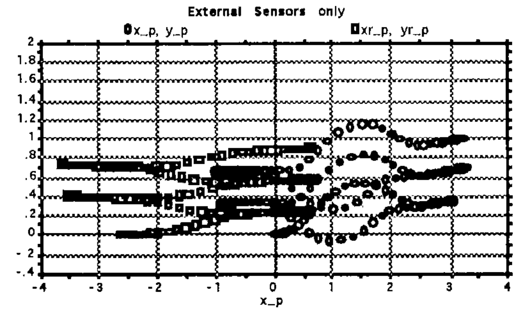

(e) External sensors

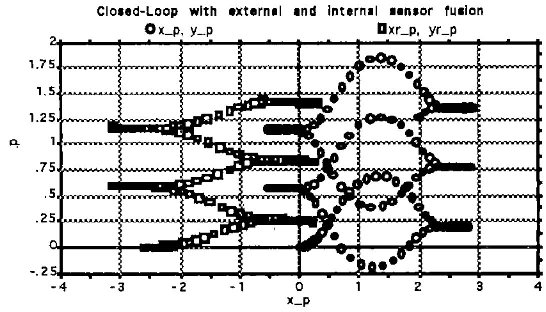

(f) Sensor fusion

Fig. 8 Multiple-maneuvers and 2-D control regimes. In cases show the position of the front (*x_p, y_p*) and rear (*xr_p, yr_p*) of the plant sampled at uniform time intervals. The data tends to bunch up at the maneuver end-points where the plant is moving slowly. All Units are in Meters.

Closed-loop sensor fusion reduced single-maneuver overshoot by 0.60 meters over the open-loop case.

If the orientation is non-zero at the maneuver end-point, the controller will bang on the steering in an attempt to correct the orientation. The plant is moving slowly at the maneuver end-point and steering changes have less effect on the plant's orientation. Even if the orientation is zero, the steering-wheel of the plant never stays at zero. This is due to steering reversals after every three bangs. The three bangs last three hundredths of a second, each of which turn the front wheel by $(1/2)50$ radians/s$^2 \cdot (0.03)^2 = 0.0225$ radians (when the steering velocity is zero at the time of the bang). This delay is due to the two integrators needed to numerically compute the steering-wheel angle from the applied steering acceleration and to the delay in the controller's decision to bang. The controller's decision to bang is delayed by a simulator time interval (0.01 seconds) because of the next-state equations used in the simulation.

Figure 8 shows several 2-D control regimes running in simulation over a twenty second interval. Note that the open-loop control with acceleration error has so much overshoot that the scale of Fig. 8 (a) had to be altered. In addition, Fig. 8 (a) shows an upper-bound on parking performance. Figure 8 (c) shows that error-free sensing does lead to a deflection smaller than the closed-loop sensor fusion shown in Fig. 8(f). The reason is that the rate at which the plant proceeds along the Y-axis is a quadratic function of the available maneuver room. Since the sensor fusion case takes more room to maneuver (due to overshoot), it is able to proceed into the space faster.

In Fig. 8(a), there is no error in the maneuver and so none accumulates. However, in Fig. 8(b), the error accumulates and must be reduced with sensor feedback. Either external or internal sensor feedback will keep the error from accumulating.

Several experiments were performed to determine the effect of modeling error in the 10% to 25% range on the overshoot, average steering-wheel angle and car-orientation errors. It was found that the change in the average steering-wheel and car-orientation error was a few thousandths of a radian and that the overshoot is increased from 3.16 meters to 3.46 meters for a PSL-L (Parking Space Length-Length of a car) of 2.4 meters, as modeling error varies from 10% to 25%.

4. Discussion and Conclusion

We have studied a new approach to parallel parking a car. We have seen that multiple maneuvers made possible by using proportional plus derivative control, and have used the nonholonomic constraints to assist in the estimation of plant-states.

Modeling sensors, whose reliability is inversely related to availability, led to the development of a new sensor-fusion technique. This technique can perform sensor fusion using sensors with systematic error, is computationally simple enough for an embedded controller, and requires no statistical assumptions.

The technique was formulated under the premise that the plant-model error is greater than the internal-sensor error and that the internal-sensor error is greater than the external-sensor error. The external sensor overrides the internal sensors, which are reset each time that an external sensor reading is received.

The observer fuses data from the plant-model and sensors by using the most recent sensor measurements to correct the plant-model's error. This results in the observer using the plant model like a nonholonomically constrained sample-and-hold.

Our sensor fusion technique reduced the amount of overshoot over using either internal or external sensors alone. Proportional plus derivative control gave us the ability to track a fifth-order polynomial, despite bang-bang control constraints and an incomplete knowledge about the steering performance of the model.

The idea for using a reference path for the car to track is not new and there are other curves which are better if the parking criterion is not used. For example, the cubic spiral is a smoother path than the fifth-order polynomial[4]. The drawback of this curve is that it has non-zero curvature at the maneuver endpoints. The advantage is that the curve is the smoothest path for autonomous vehicles and this will minimize the jerk exerted on passengers.

The work on configuration-space search[5] provided a theoretical basis for the nonholonomic constraints used in this paper. The draw-back of configuration-space search is that it requires considerable computational resource. With the introduction of a simple parking criterion, the selection of a path becomes a computationally tractable problem. This approach to reference path generation eliminates the heavy computational machinery needed for configuration-space search.

The advantage of using a parking criterion goes beyond the elimination of heavy computational machinery. For example, fuzzy control and neural network control run the risk of being suboptimal with respect to the parking criterion. With these approaches, the operator's skill becomes the limiting factor. In addition, the use of the parking criterion has led to fore-knowledge of the time it will take to park. This information is not available with the fuzzy or neural network control approaches. Literature shows that the neural-network approach requires thousands of training sessions and this is not needed with the approach used in this paper[6].

This is the first work known to the author to combine nonholonomic constraints and procedural knowledge to perform sensor-fusion. It should be noted, however, that the idea of sensor fusion in the robotic vehicle is not new[7], but the assumption of systematic error and sensor fusion with sensors whose reliability is inversely related to their availability is novel.

Literature reveals that the fuzzy controller based maneuvers can result in collision when faced with modeling errors[8]. With the approach used in this paper, collision is eliminated (given an *a priori* knowledge of maximum modeling and sensor error).

When braking performance is overestimated, it is impossible to completely eliminate overshoot and so we require a buffer zone around the car. For the 2.4 meter maneuver, the plant overshot by 0.8 meters, assuming 25% acceleration and deceleration error.

Each experiment varied a single parameter and several results of interest were found. For example, the plant appeared to be stable with steering accelerations of 25 radians/s^2. The bang-bang acceleration on the steering produced a finite number of resulting steering-wheel angles (8 for 50 radians/s^2 and 14 for

25 radians/s^2). This is due to small oscillation of the path about the reference curve, which results in steering reversals after every few bangs. The bangs vary from three to five hundredths of a second, each of which turn the steering wheel by $(1/2)\cdot 25$ radians/$s^2 \cdot 0.05 = 0.031\,25$ radians (if the steering velocity were zero at the time of the bang). If steering bangs lasted longer, the steering wheel of the test vehicle could be turned from one steering extreme to the other in about three tenths of a second. We have also found that reduction of steering acceleration can cause tracking overshoot which results in an increase in the maximum steering speed. Bang-bang control of a car is not new and is cited as a method for moving a car in minimum time[9]. However, bang-bang control on the steering system of a car is novel and so there is no literature available with which we may compare our simulation results.

Proportional plus derivative control is well established[10] and when the reference curve is known *a priori* such control is not optimal with respect to tracking error[11]. However, when the steering system is subject to bang-bang control, proportional plus derivative control was the simplest control law to work with reasonable tracking error.

The stability of a class of non-linear steered systems is discussed for ship models in Ref.(12). The question of how to apply stability analysis to bang-bang control laws on systems with nonholonomic constraints remains open.

References

(1) Lyon, D., Parallel Parking with Nonholonomic Constraints, Ph.D. Thesis, Computer and Systems Engineering, Rensselaer Polytechnic Institute, Troy, NY, 12181, (1991)

(2) Parallel Parking a Car with Nonholonomic Constraints, IEEE Intelligent Vehicles, Detroit, MI, (1992).

(3) Ad-Hoc and Derived Parking Curves, SPIE-International Society for Optical Engineering, Boston, MA, (1990).

(4) Kanayama, Y. and Hartman, B., Smooth Local Path Planning for Autonomous Vehicles, TRCS 88-15, (1988), Department of Computer Science, University of California, Santa Barbara, CA.

(5) Barraquand, J. and Latombe, J., Robot Motion Planning : A DistributedRepresentation Approach, STAN-CS-89-1257, (1989), Department of Computer Science, Stanford University Stanford, CA, 94305.

(6) Widrow, B. and Nguyen, D., Improving the Learning Speed of 2-Layer Neural Networks by Choosing Initial Values of the Adaptive Weights, Int. Joint Conf. on Neural Networks, (1990), pp. 21-26.

(7) Maybeck, P., The Kalman Filter : An Introduction to Concepts, Autonomous Robot Vehicles, Editors Cox, I. and Wilfong, G., (1990), Springer-Verlag, NY.

(8) Sugeno, M. and Murakami, K., An Experimental Study on Fuzzy Parking Control Using a Model Car, Industrial Applications of Fuzzy Control, (1985), North Holland Publishing, NY.

(9) Kirk, D., Optimal Control Theory, (1970), Prentice-Hall Inc., Englewood Cliffs, NJ.

(10) Frederick, D. and Carlson, B., Linear Systems in Communication and Control, (1971), John Wiley and Sons, Inc., NY.

(11) Athans, M. and Falb, P., Optimal Control : An Introduction to the Theory and Its Applications, (1966), McGraw-Hill Book Company, NY.

(12) Fossen, T.I. and Lauvdal, T., Nonlinear Stability Analysis of Ship Autopilots in Sway, Roll and Yaw, In Proceedings of the 3rd Conference on Marine Craft Maneuvering and Control (MCMC '94), (1994), pp. 113-124, Southampton, UK, http://citeseer.nj.nec.com/fossen94nonlinear.html