System Identification and Robust Control of Farm Vehicles Using CDGPS

Gabriel Elkaim, Michael O'Connor, Thomas Bell, and Dr. Bradford Parkinson, *Stanford University*

BIOGRAPHY

Gabriel H. Elkaim is a Ph.D. candidate in Aeronautics and Astronautics at Stanford University. He received his B.S.E. in Aerospace Engineering from Princeton University in 1990 and his M.S. from Stanford in 1995. He has worked for Schlumberger Wireline Logging and Testing in Montrouge, France.

Michael L. O'Connor is a Ph.D. candidate in Aeronautics and Astronautics at Stanford University. He received his B.S. in Aeronautics and Astronautics from MIT in 1992 and his M.S. from Stanford in 1993. He is expected to receive his Ph.D. in December 1997.

Thomas Bell is a Ph.D. candidate in Aeronautics and Astronautics at Stanford University. He received his B.S. in Mechanical Engineering from Cornell University in 1991. He is a former communications officer in the United States Army.

Bradford W. Parkinson, Ph.D., is professor of Aeronautics and Astronautics at Stanford University, and Program Manager of the Relativity Gyroscope Experiment (Gravity Probe B). He served for six years as the first Program Director of the GPS Joint Program Office, and has been instrumental in GPS program development. Dr. Parkinson heads the NASA Advisory Council and is a Fellow of the AIAA and a member of the National Academy of Engineering.

ABSTRACT

Automatic control of agricultural vehicles has been a research goal for many years. Previous attempts have failed largely due to sensor limitations. With the advent of modern GPS receivers, a single low-cost sensor has been synthesized in which centimeter-level position and attitude measurements of the vehicle state are available using Carrier-Phase Differential Global Positioning System (CDGPS). Previous research at Stanford University has demonstrated control of a Deere 7800 tractor at low speed without implements.

Utilizing the Observer/Kalman Filter Identification method, input-output measurements were combined to form a central model of the farm tractor for all speed and implement combinations. A linear quadratic regulator (LQR) controller was designed utilizing the central model and experiments were performed to measure the achieved performance.

Tight line-following was achieved that surpasses the capability of human drivers—at a speed of 1.75 m/s (4 mph), standard deviations for tractor with any implement were less than 7 cm. (3 in.) and with no implement, the standard deviation at this speed is less than 5 cm. (2 in.). Tractor automatic control performance improves as tractor velocity decreases, but remains excellent within the usable range—standard deviations remain at less than 11 cm. (4 in.) for all implement and speed combinations.

INTRODUCTION

The primary goal of this work was to experimentally demonstrate system identification and precision closed-loop control of a farm tractor using CDGPS as the only sensor of vehicle position and attitude.

Autonomous guidance of ground vehicles has been a research objective for many years. Previous attempts have failed largely due to sensor limitations—experimental systems required cumbersome auxiliary equipment in or around the field [1,2], while others have relied on vision systems that require clear daylight, good weather, or field markers that require complex pattern recognition [3,4]. With the advent of modern GPS receivers, a single, low-cost sensor for measuring multiple vehicle states has been incorporated into a prototype self-guided tractor.

Previous work done at Stanford demonstrated carrier-phase differential GPS (CDGPS) as the primary state

Figure 1 — GPS-Equipped Tractor

sensor, with centimeter-level positioning accuracy [5] and attitude determination to better than 0.1° [6]. Using CDGPS as the vehicle state sensor, and augmenting the GPS satellite signals with a ground-based "pseudo"-satellite, integer cycle ambiguity resolution is assured, and the solution integrity is greatly enhanced [7]. These extremely accurate and reliable measurements of multiple vehicle states lend themselves to system identification, estimation, and automatic control [8].

Previous work at Stanford has concentrated on setting up the mechanisms to achieve automatic control of a large agricultural vehicle, initially working with a golf cart [9], and progressing to a large farm tractor [10]. Automatic control of ground vehicles at Stanford has progressed in measured steps. In 1995, a golf cart achieved line-following using CDGPS to a standard deviation of 5 cm. (2 in.)[9]. In 1996, a simple kinematic model was used to demonstrate tractor control about a line to a standard deviation of 2.45 cm. (1 in.) at a velocity of 0.33 m/s (0.75 mph)[10]. At the ION-GPS 1996 conference, tractor control was demonstrated using identified models from input-output data to control the tractor about a straight line with a standard deviation of less than 2.5 cm. (1 in.) at 0.33 m/s (0.75 mph) and less 7.2 cm. (3 in.) at a speed of 1.6 m/s (3.5 mph)[8].

Continuing the work presented in 1996, in early 1997 the Stanford group demonstrated integer cycle ambiguity resolution using a curving trajectory about a single pseudolite [11]. The basic shortcoming of this research was that while precise control was demonstrated and proved very effective, each controller was designed to operate at one specific speed, and with one kind of implement. Switching between controllers had to be accomplished via driver interaction, gain scheduling, or additional sensors that allowed the software to determine the configuration and speed of the tractor and switch between these individual controllers.

A robust control system was required that utilized a single controller capable of controlling the tractor at all speeds with any implement without user interaction. This was a challenging task—fixed gain controllers that are designed for a specific speed yield poor performance at speeds slower than the design point and will go unstable at speeds greater than the design speed. Furthermore, different implements can greatly alter the dynamics of the tractor to the point that controller instability results. Many non-linear effects come into play when an implement is towed through the soil: friction on the implement will vary with soil conditions, soil density variations will impart lateral forces on a running implement, and strange tractor/implement interactions such as implement "hop" can occur at certain speeds.

"Hop" manifests itself when a towed implement causes the tractor to bounce up and down over the soil as the "spring" of the rubber tire loads and unloads because of the towing force. Obviously, during these conditions, front wheel contact is time-varying and directly affects the tractor's response to any steering commands.

All implements tend to reduce the effect of front steer angle on the tractor's trajectory due to loss of contact or reduced normal load on the ground. Disturbances are much larger and, in many cases, both time-varying and non-linear. In general, hitched implements (ie: rigidly attached to the tractor) transmit lateral disturbances from the soil directly to the tractor, whereas towed implements do not. Towed implements can drag the tractor to one side or the other due to non-symmetrical loading of the tractor. For instance, the disker has angled wheels that tend to pull the implement to the left.

This paper explores the effort to develop one controller that can address all of these issues, demonstrates the level of control that was attained with this "central" controller, and compares the results to previous designs as well as an expert human driver.

THE HARDWARE SETUP

Vehicle Hardware: The test platform used for vehicle identification and control testing was a John Deere Model 7800 tractor (Fig. 1). Four single-frequency GPS antennas were mounted on the top of the cab, and an equipment rack was installed inside. Front-wheel angle was sensed with a potentiometer—the only non-GPS sensor used in the system—and actuated using a modified Orthman electro-hydraulic steering unit. A Motorola MC68HC11 microprocessor board was the communications interface between the control computer and the steering unit. The microprocessor converted serial commands from the control computer into pulse-width modulated signals

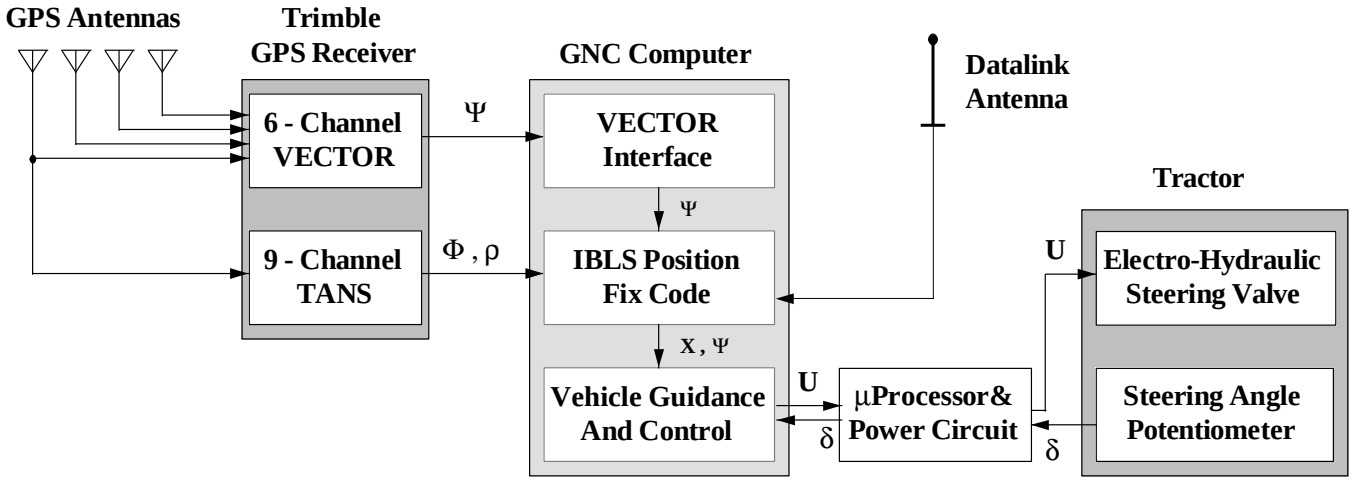

Figure 2 — GPS Hardware Architecture

which were sent to the power circuits that controlled the steering valves. Wheel position was sampled and digitized by the MC68HC11 and sent to the controls computer at 20 Hz. Both the steering unit and the wheel angle potentiometer were highly non-linear devices.

GPS Hardware: The equipment rack within the tractor housed the CDGPS-based system used for vehicle position—this system was similar to the one used by the Integrity Beacon Landing System (IBLS)[5]. A four-antenna, six-channel Trimble Vector receiver produced attitude measurements at 10 Hz. A single-antenna Trimble TANS PC-card receiver produced code- and carrier-phase measurements used to calculate vehicle position at 5 Hz. An Industrial Computer Source Pentium-based PC running the LYNX-OS real-time operating system performed the attitude interface, position calculations, data collection and controls calculations using software written at Stanford. A functional representation of the tractor hardware appears below in Figure 2.

Reference Station: The ground reference station that produced the carrier-phase differential corrections consisted of an Industrial Computer Source Pentium PC, a TANS 9-channel PC-card receiver, and software written at Stanford to broadcast phase corrections to the tractor through a Pacific Crest radio modem at 4800 bps.

Pseudolite: CDGPS required a method for initializing the solution to the correct position. A robust method for integer-cycle-ambiguity resolution was to travel a curved trajectory about a single pseudolite, and use the geometrical leverage to pinpoint the exact solution [11]. The pseudolite was an IntegriNautics L1-pseudolite transmitter mounted on top of a 6.7 m. (22 ft.) tall aluminum mast. A Micropulse quadriphilar helix antenna is used without a ground-plane in order to radiate power below the local horizon in the direction of the tractor's patch antennae.

VEHICLE MODELING AND SYSTEM IDENTIFICATION DATA COLLECTION

Agricultural farm vehicles must be able to operate over various types of terrain and with a variety of implements. While previous work at Stanford has demonstrated closed-loop line following based on a simple kinematic model (Figure 3) to a remarkable precision [10], the model is based on assumptions that are known to be false.

The kinematic model, based on simple geometry rather than inertias and forces, assumes both a constant velocity

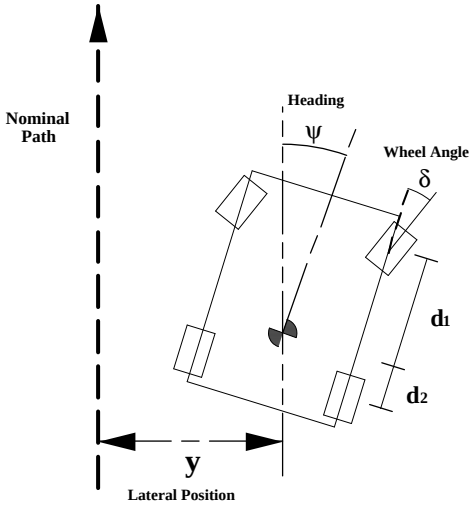

$$\begin{bmatrix} \dot{y} \\ \dot{\psi} \\ \dot{\delta} \end{bmatrix} = \begin{bmatrix} 0 & V & \frac{Vd_2}{(d_1+d_2)} \\ 0 & 0 & \frac{V}{(d_1+d_2)} \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} y \\ \psi \\ \delta \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} u$$

Figure 3 — Simple Kinematic Model

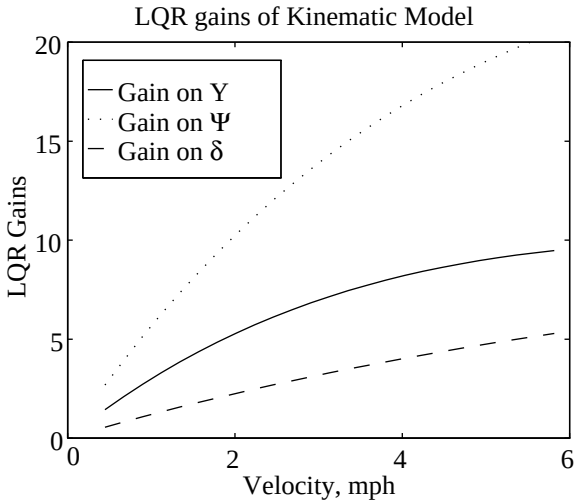

Figure 4 — LQR gains vs. Velocity

along the path as well as no lateral wheel slip. While the velocity may not vary a great deal, it is not constant, and the four-wheel-drive on the tractor cannot move the vehicle forward without slipping the wheels.

The kinematic model did, however, provide an initial reference point to examine the main parameters that affect tractor dynamics. Close inspection of the kinematic model showed, as observed by experimentation, tractor dynamics are a strong function of forward velocity. Mathematically, both the cross-track deviation (y) and the azimuthal deviation (Ψ) integrate not with time, but rather with distance travelled. The steering angle (δ), however, is purely the time integral of the control signal. In order to extract the relevant time to distance transformations, an ideal linear quadratic regulator (LQR) controller was designed for each velocity using the kinematic model. The results are plotted in Figure 4—clearly the feedback gains on crosstrack and azimuth are parabolic functions of velocity, but the gain on steering angle remains linearly proportional to velocity throughout the speed range.

In light of these results, new outputs to formulate velocity-invariant controllers were synthesized from crosstrack and azimuth normalized by velocity, and wheel angle as

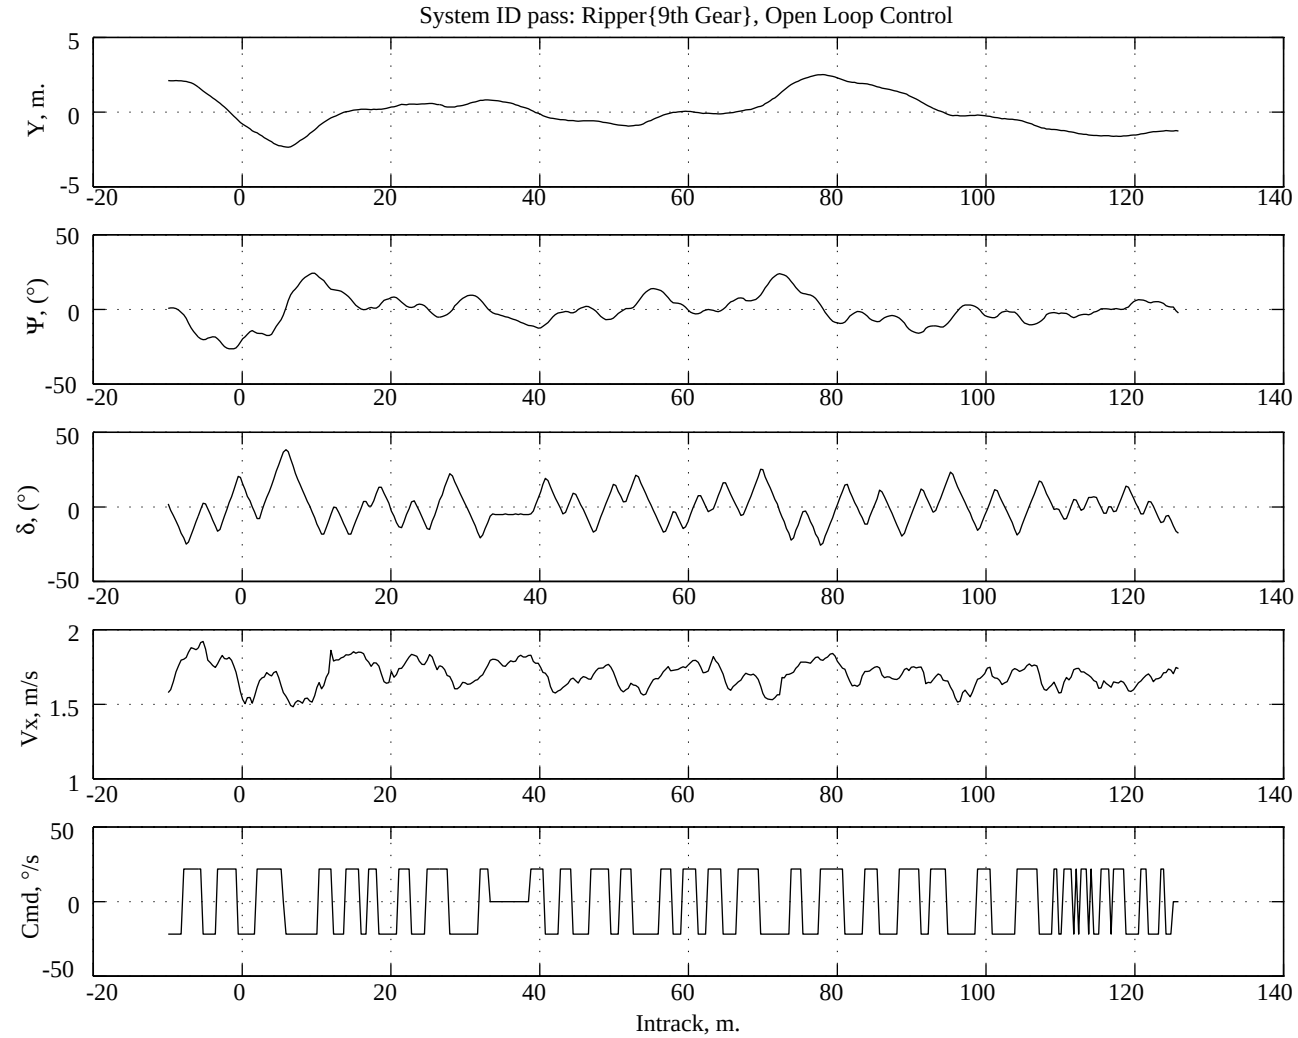

Figure 5 — Typical System Identification Pass

measured. In order to gather data to perform a proper system identification of the tractor, a series of open-loop line-following tests were conducted in which a human driver, through the GNC computer, caused the steering to either slew left or right at the maximum steering rate. Also, the driver commanded the steering rate to zero through the electro-hydraulic actuator in order to track a roughly straight line. This "pseudo"-random input was designed to apply the maximum power to the tractor through the controls and produce a rich output that would contain information from all modes of interest. A typical pass for system identification is pictured on the previous page in Figure 5.

These data passes were run with the tractor unencumbered, with a towed disker, and a hitched three-shank ripper. Data was collected at velocities of 1.0, 1.25, 1.75, 2.25, 3.0 m/s (corresponding to 5th, 7th, 9th, 11th and 13th gears) and was subsequently used for calculation and validation of a linear plant model. The data was gathered separately and then post-processed as a "batch" for identification. The data was pre-processed and used to identify local controllers for each speed and implement combination as well as velocity-invariant controllers for each implement. Using the aggregate of all open-loop identification data, a central model that accurately predicted the behavior of the tractor under all conditions was identified and a corresponding controller designed with this model.

The controller was designed using a standard LQR methodology. The quadratic cost, as calculated below, minimized the weighted sum of the outputs (y_{max} and u_{max} are design parameters).

$$J = \sum_{k=0}^{\infty} \left(\vec{x}_k^T C^T \begin{bmatrix} \frac{1}{y_{\max}^2} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} C\vec{x}_k + \vec{u}_k^T \left[\frac{1}{u_{\max}^2} \right] \vec{u}_k \right)$$

THE OBSERVER/KALMAN IDENTIFICATION PROCESS

The method of identifying the plant models, as well as the process and sensor noise statistics, chosen for this project was the observer/Kalman filter identification (OKID) method. This method of system identification uses only input and output data to construct a discrete-time state-space realization of the system. Since OKID's development at NASA Langley for the identification of lightly-damped space-structures, many advances on the basic theory have been published[13]. Given a linear discrete-time state-space system, the equations of motion can be written as follows:

$$x_{k+1} = Ax_k + Bu_k$$
$$y_k = Cx_k + Du_k$$

It has been shown that the triplet, [A,B,C] is not unique, but can be transformed through any similarity transform (ie. the outputs are unique, but the internal states are not). However, the system response from rest when perturbed by a unit pulse input, known as the system Markov parameters, are invariant under similarity transforms. These Markov parameters are:

$$Y_0 = D,\ Y_1 = CB,\ Y_2 = CAB,\ \ldots,\ Y_k = CA^{k-1}B$$

When these Markov parameters are assembled into a specific form—the generalized Hankel matrix—this matrix can be decomposed into the Observability matrix, a state transition matrix, and the Controllability matrix; thus the Hankel matrix (in a noise-free case) will always have rank n, where n is the system order.

$$H(k-1) = \begin{bmatrix} Y_k & Y_{k+1} & \cdots & Y_{k+\beta-1} \\ Y_{k+1} & Y_{k+2} & \cdots & Y_{k+\beta} \\ \vdots & \vdots & \ddots & \vdots \\ Y_{k+\alpha-1} & Y_{k+\alpha} & \cdots & Y_{k+\alpha+\beta-1} \end{bmatrix}$$

$$H(k-1) = \begin{bmatrix} C \\ CA \\ CA^2 \\ \vdots \\ CA^{\alpha-1} \end{bmatrix} A^{k-1} \begin{bmatrix} B & AB & A^2B & \cdots & A^{\beta-1}B \end{bmatrix}$$

Because noise will corrupt this rank deficiency of the Hankel matrix (the Hankel matrix must always be full rank) the Hankel matrix is truncated using a singular value decomposition (SVD) at an order that sufficiently describes the system. This truncated Hankel matrix is then used to reconstruct the triplet [A,B,C] in a balanced realization that ensures that the controllability and observability Grammians are equal. This is referred to as the Eigensystem Realization Algorithm (ERA); a modified version of this algorithm that includes data correlation is used to identify the tractor. A more complete treatment of the subject can be found in [13].

For any real system, however, system pulse response cannot be obtained by simply perturbing the system with a pulse input. A pulse with enough power to excite all modes above the noise floor would likely saturate the actuator or respond in a non-linear fashion. The pulse response of the system can, however, be reconstructed from a continuous stream of rich system input and output behavior. Under normal circumstances, there are not enough equations available to solve for all of the Markov parameters. Were the system asymptotically stable, such that A_k=0 for some

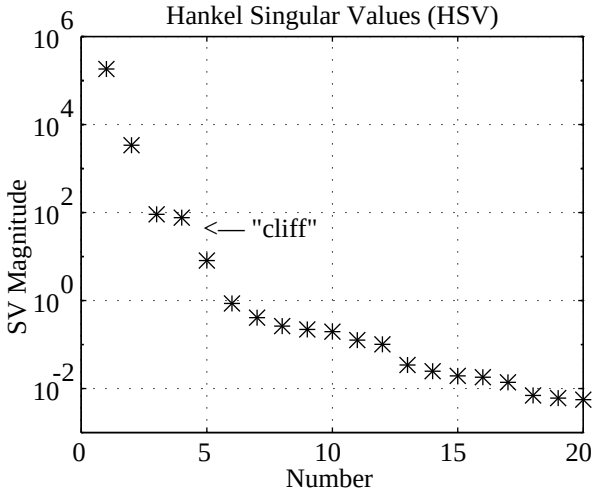

Figure 6—Hankel SVD forAgregate Tractor Data

k, then the number of unknowns could be reduced. The identification process would be of little value if it could only work with asymptotically stable systems.

By adding an observer to the linear system equations, the following transformation can take place:

$$x_{k+1} = Ax_k + Bu_k + Gy_k - Gy_k \quad \text{[add zero]}$$
$$x_{k+1} = [A + GC]x_k + [B + GD]u_k - Gy_k$$
$$x_{k+1} = \hat{A}x_k + \hat{B}v_k$$

$$\hat{A} \equiv [A + GC],\ \hat{B} \equiv [B + GD \quad -G],\ \text{and } v_k \equiv \begin{bmatrix} u_k \\ y_k \end{bmatrix}$$

Thus, the system stability can be augmented through an observer and the ideal Markov parameters established through a least-squares solution [14]. It is useful to note that the realization also provides a pseudo-Kalman observer. The observer orthogonalizes the residuals to

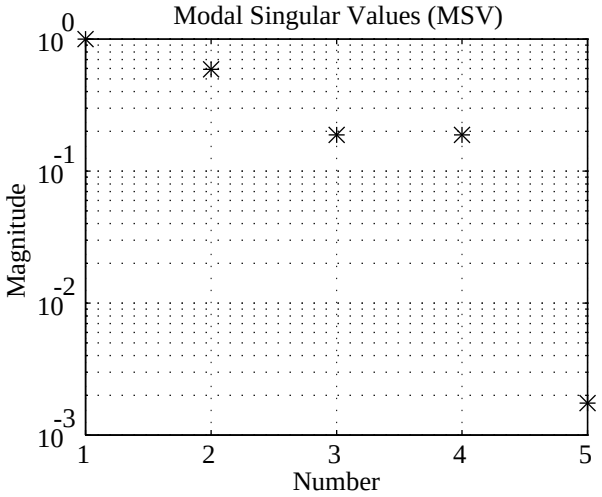

Figure 7—Modal SVD forAgregate Tractor Data

time-shifted versions of both input and output. Utilizing the separation lemma and the provided Kalman filter, only the controller gains need be designed to implement a full-state-feedback linear quadratic gaussian (LQG) controller. An improved version of the OKID process, which includes residual whitening [15], was used to identify the farm tractor from the experimental data.

In order to compare velocity invariant and velocity-specific controllers, the data runs for each velocity and implement combination were combined and identified to generate plant models and estimators that were then used to create local controllers for each condition. In addition, each implement had all of the data runs at all of the velocities with the crosstrack and azimuth measurements normalized by velocity combined; the resulting sets were again identified and used to generate velocity-invariant-controllers for no implement, towed implement, and a hitched implement.

Lastly, all of the open-loop data was combined and identified to create a central model—one model that predicted all variations to both velocity and implement differences. This model was used to generate one central controller capable of stabilizing the tractor at all speeds and with every tested implement. The identification of all of the open-loop data produced a model that is essentially a best fit in a least squares sense to all possible combinations of velocity and implement.

An SVD of the aggregate velocity-normalized data for the tractor demonstrated a large drop in the magnitude of the singular values from the fourth to the fifth, indicating a system of order four, n=4 (Figure 6). In addition, modal singular values (Figure 7) of all tractor models of order higher than four exhibited a two order-of-magnitude drop from the fourth modes to modes higher than four.

As an experimental check, on the following page in Figure 8, the identified model (with observer) was compared to a typical system identification pass. The match was excellent between the modeled and actual data. Note that the reconstructed data is offset in the figure to distinguish between the actual and reconstructed data. The matched data was normalized to velocity, thus the units are seconds (meters / meters per second) for crosstrack, and degrees-seconds per meter for azimuth, and is summarized in Table 1 below.

	bias	1–σ
Lateral, y/Vx (s)	0.001	0.0502
Azimuth, Ψ/Vx (°s/m)	0.041	0.764
Steering, δ (°)	0.148	2.407

Table 1—Reconstruction vs. Actual Data

EXPERIMENTAL RESULTS

In order to validate this approach to velocity- and

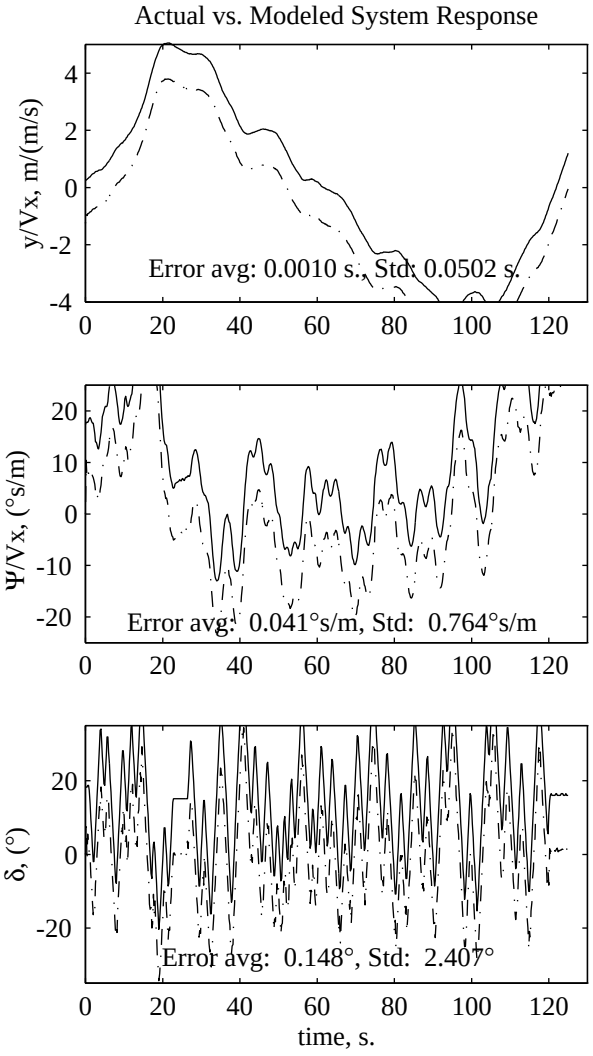

Figure 8 — OKID Reconstruction vs. Actual Data

implement-invariance, a series of experiments were run and the results collected. In order to establish the boundary of how well each velocity-invariant controller would be expected to perform, the experiments were performed with each fixed-velocity controller (ie. the 5th gear subsoiler controller was designed using data identified solely from open-loop data taken at 5th gear with the subsoiler in the ground).

Next, a series of passes were run down the test field with each velocity-invariant controller, at every speed and with that controller's designed implement (ie. all of the subsoiler data was used to create a velocity-invariant subsoiler controller). Lastly, the central controller was tested at all speeds and with the three different implement combinations. Table 2, on the following page, tabulates the central controller results—the units are meters for crosstrack biases and standard deviations. As expected, the biases and standard deviations increased with velocity, but were well within the requirements of the agricultural industry.

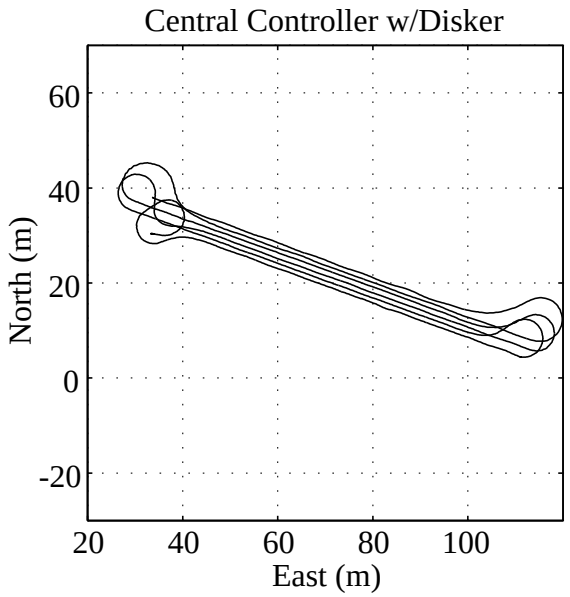

Figure 9 — Overhead View of Tractor Path

Figure 9 shows an overhead view of a typical controlled pass with several rows, each one meter from the other, that were disked at various speeds. Figure 10 shows graphically the standard deviations for the various controllers at the different speeds—three velocity-invariant controllers are compared with each local controller, specific to their own implement; these are the velocity-invariant disker, subsoiler, and unencumbered controllers.

In order to compare the computer controlled results to a human driver, tests were run at a farm with an expert human driver. A hitched 10-row bedder was attached to the

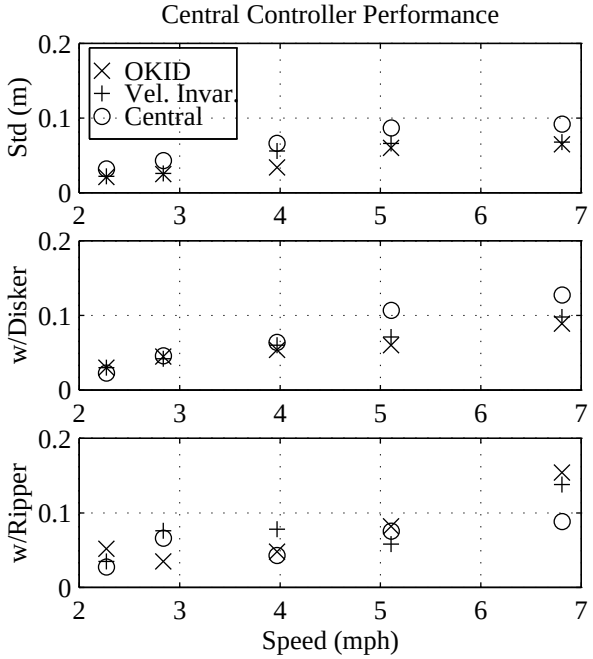

Figure 10—Standard Deviations vs. Speed

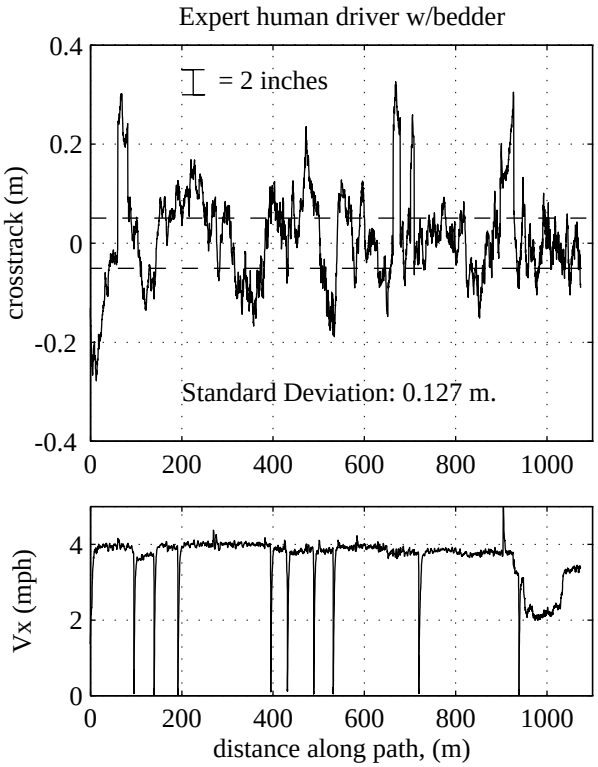

Figure 11—Expert Driver with Hitched Implement

tractor and the driver was instructed to drive as straight a line a possible. The data from this experiment was recorded and the best line was fitted through the data. Deviations from this line were calculated and appear in Figure 11. Note that this is an exceptional driver driving under ideal conditions. This driver was able to keep the tractor on the line with a standard deviation of 12.7 cm. (5 inches). Velocity drops several times when the driver stopped to let dust blow by the cab in order to maintain visibility.

Figure 12, on the next page, is a close-up view of the central controller at 1.75 (4.0) and 3.0 m/s (6.7 mph). The graphs show the overlay of several passes with disker and subsoiler automatically guided. All of the controllers guide the tractor precisely down the desired path with or without an implement, even at speeds that were previously impossible while maintaining stability.

An interesting comparison is the crosstrack error of the central controller at 1.75 m/s (4.0 mph) to the crosstrack error of the expert human driver. Note that the vertical scales of the two plots are identical, and the bands that mark two inches off of the desired track are also identical. It is entirely clear that the GPS guidance yields much more precise control than the human driver.

The central controller is able to control the hitched implement at the same speed to a standard deviation of less than 7 cm. (3 in.), nearly twice the precision of an expert human driver. Using automatic guidance, the tractor did not need to stop for lack of visibility. Close examination of Figure 12 shows that the central controller shows no clear “trends” in the errors, unlike the human driver. The human driver was given the benefit of choosing his own line across the field, defined mathematically as a least-squares fit to the data. In side-by-side comparisons, it is expected that the automatic control will show further improvement over human driving. Clearly, there are far-reaching implications for this kind of system.

AUTOMATIC CONTROL IMPLICATIONS

With a demonstrated ability to both align the tractor to the start of any given row and the ability to precisely guide the tractor down that row, many previously difficult tasks become possible. By controlling the tractor better than an expert human driver or by matching that driver’s accuracy at a higher speed, automatic control of farm vehicles can save time, increase farm efficiency, and lower overall costs.

As an example, in Arizona, where water is scarce, farmers are using tape irrigation, a technique which deposits water directly on the plant roots. A rough figure that is quoted for maintenance costs due to severed tapes is approximately $100,000 per year on a large farm. This cost could be largely mitigated with automatic control through reduced damage to the tapes from implements as well as placing the plants close to the irrigation tape for optimum growth. In this environment, less water would be required and productivity would increase while lowering costs.

Another consideration is that, for non-precision applications, such as disking and ripping, automatic control would enable the same level of precision as a human driver at higher speed and without driver fatigue. Current implements are not designed to run at speeds of greater than 2 m/s (4-5 mph) because the bearings on the disker wear too quickly and the load on the subsoiler shanks at these speeds becomes excessive. New implements could be designed to take advantage of the higher speeds enabled by automatic control.

	Unencumbered		Disker		Ripper	
	bias (m)	1-σ (m)	bias(m)	1-σ (m)	bias(m)	1-σ(m)
5th gear (2.25 mph)	0.009	0.032	0.004	0.023	0.004	0.027
7th gear (3.0 mph)	-0.002	0.043	0.007	0.046	0.007	0.066
9th gear (4.0 mph)	-0.001	0.066	0.035	0.064	-0.010	0.043
11th gear (5.0 mph)	0.007	0.087	0.060	0.107	0.012	0.075
13th gear (6.5 mph)	-0.001	0.092	0.085	0.127	0.075	0.088

Table 2—Central Controller Experimental Results

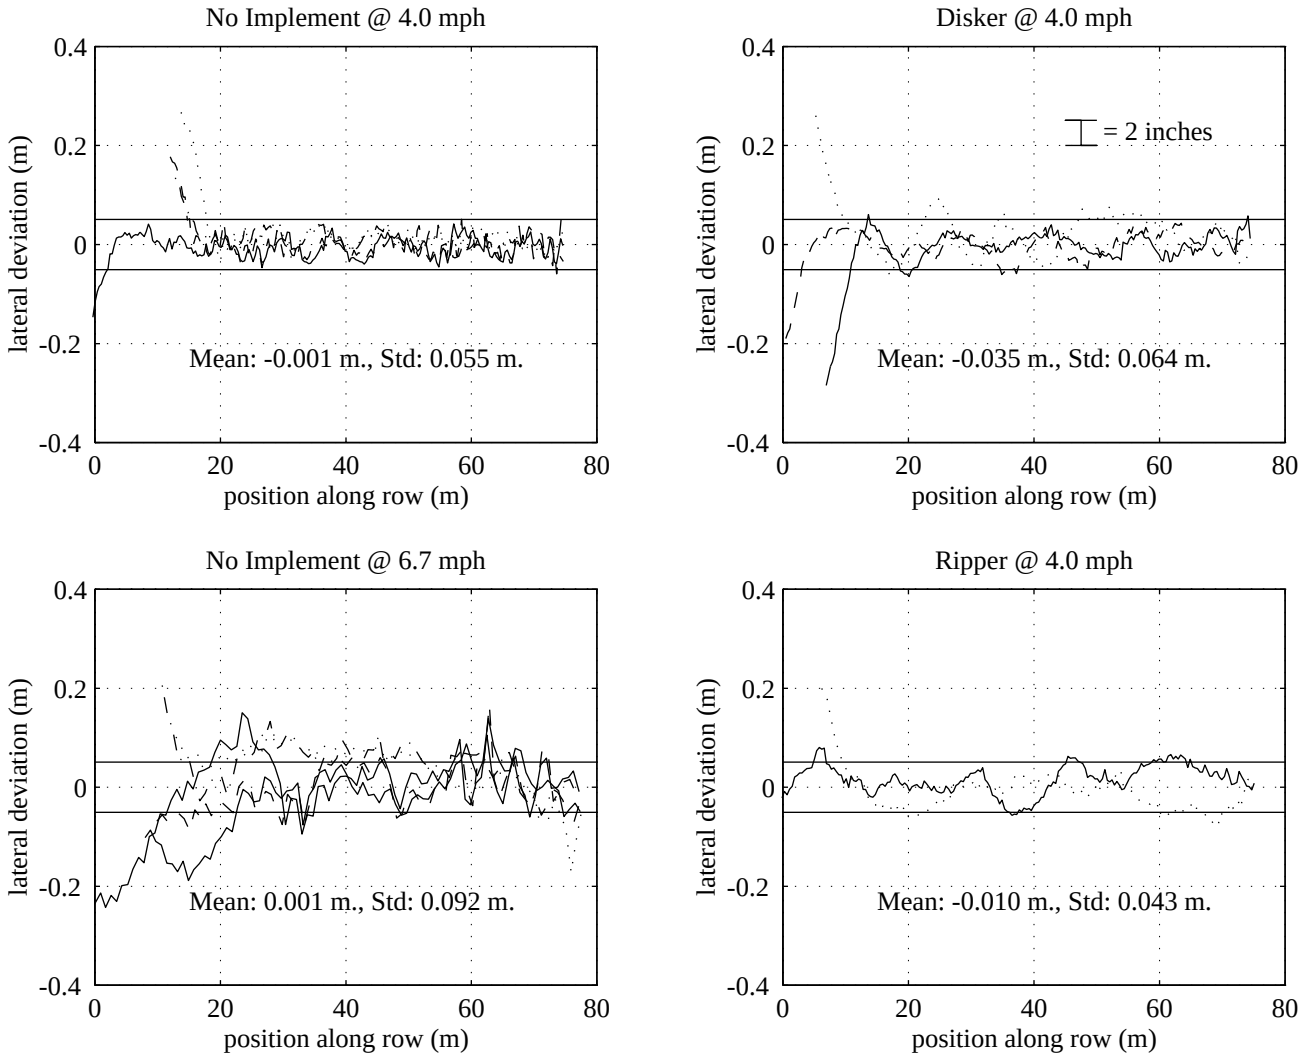

Figure 12 — Close-in View of Tractor Passes

CONCLUSIONS

It has been demonstrated that, pursuing the research presented in 1996 [8], controllers can be designed using only input-output data streams that stabilize and drastically reduce the error in straight-line-following applications. The controllers have shown remarkable performance even when operating with towed or hitched implements over uneven ground.

Finally, one controller was demonstrated that was both velocity- and implement-invariant which stabilized the tractor in all modes of operation—hitched, towed, or no implement, and at any speed. This is a very necessary step in order to transition from a laboratory experiment into a viable commercial product.

Line-following accuracies have been demonstrated that are less than 11 cm. (4 in.) standard deviation at all speeds with all implements. When tractor velocity is reduced to normal operating range, the 1-σ values for line-following drop to less than 5 cm. (2 in.) even when towing an implement through a rough field.

At this juncture, the requirements of most farms are met with the existing system. To make the inner workings of the system transparent to the user, automatic controllers were implemented that are both velocity- and implement-invariant. Tight-line following will be accomplished with no further intervention by the driver. Using the existing system, extraordinary accuracy can be attained without the requirement for clear vision and without fatiguing the driver.

Greater speeds are enabled by automatic control that can contribute to increased farm productivity. Advances in irrigation combined with the ability for the tractor to know the exact location of the implement using CDGPS enable a wide variety of benefits including a reduction in water

consumption and an increase in productivity.

Utilizing advanced time-domain system identification techniques, design of a high-gain central controller that is both velocity- and implement-invariant was accomplished with only input-output measurements from CDGPS. Robust automatic control at various speeds and with several test implements was demonstrated experimentally, achieving standard deviations of less than 7 cm. (3 in.) at all useable speeds.

ACKNOWLEDGMENTS

The authors would like to thank several groups and individuals who made this research possible. At Stanford, the entire GPS group has been extremely helpful, with Boris Pervan, Stu Cobb, and Dave Lawrence deserving exceptional mention. We would also like to thank those members of the Tractor group that are not co-authors on this paper: Ajit, Andy, Arti, Seebany, and VK Jones. Trimble Navigation provided the GPS equipment used to conduct the experiments. Funding for this research was provided by the FAA and by Deere and Company.

REFERENCES

1. Young, S.C., Johnson, C. E., and Schafer, R. L. *A Vehicle Guidance Controller*, Transactions of the American Society of Agricultural Engineers, Vol. 26, No. 5, 1983, pp. 1340-1345.
2. Palmer, R. J. *Test Results of a Precise, Short Range, RF Navigational / Positional System*, First Vehicle Navigation and Information Systems Conference - VNIS '89, Toronto, Ont., Canada, Sept. 1989, pp 151-155.
3. Brown, N. H., Wood, H. C., and Wilson, J. N. *Image Analysis for Vision-Based Agricultural Vehicle Guidance*, Optics in Agriculture, Vol. 1379, 1990, pp. 54-68.
4. Brandon, J. R., and Searcy, S. W. *Vision Assisted Tractor Guidance for Agricultural Vehicles*, International Off-Highway and Powerplant Congress and Exposition, Milwaukee, WI, Sept. 1992. Publ by SAE, Warrendale, PA, pp. 1-17.
5. Cohen, C. E., et al. *Autolanding a 737 Using GPS Integrity Beacons*, Navigation, Vol. 42, No. 3, Fall 1995, pp 467-486.
6. Cohen, C. E., Parkinson, B. W., and McNally, B. D., *Flight Tests of Attitude Determination Using GPS Compared Against an Inertial Navigation Unit, Navigation*, Vol. 41, No. 1, Spring 1994, pp 83-97.
7. Pervan, B. S., Cohen, C. E., and Parkinson, B.W. *Integrity Monitoring for Precision Approach Using Kinematic GPS and a Ground-Based Pseudolite*, Navigation, Vol. 41, No. 2, Summer 1994, pp 159-174.
8. Elkaim, G.H., O'Connor, M.L., Bell, T., and Parkinson, B. W. *System Identification of a Farm Vehicle Using Carrier-Phase Differential GPS*, Proceedings of ION GPS-96, Kansas City, MO, Sept. 1996, pp 485-494.
9. O'Connor, M.L., Elkaim, G.H., and Parkinson, B. W. *Kinematic GPS for Closed-Loop Control of Farm and Construction Vehicles*, Proceedings of ION GPS-95, Palm Springs, CA, Sept. 1995, pp 1261-1268.
10. O'Connor, M.L., Bell, T., Elkaim, G.H., and Parkinson, B.W., *Automatic Steering of Farm Vehicles Using GPS*, 3rd International Conference on Precision Farming, Minneapolis, MN, June 1996,
11. O'Connor, Bell, Elkaim, Parkinson, *Real-Time CDGPS Initialization for Land Vehicles Using a Single Pseudolite*, Proceedings of the Institute of Navigation National Technical Meeting, Santa Monica, CA,, January 14-16, 1997.
12. Cobb, H.S., Lawrence, D.G., Pervan, B.S., Cohel, C.E., Powell, J.D., and Parkinson, B. W. *Precision Landing Tests with Improved Integrity Beacon Pseudolites*, Proceedings of ION GPS-95, Palm Springs, CA, Sept. 1995, pp 1361-1368.
13. Juang J.-N. Cooper, J.E., and Wright, J.R., *An Eigensystem Realization Algorythm Using Data Correlations (ERA/DC) for Modal Parameter Identification*, Control-Theory and Advanced Technology, Vol. 4, No. 1, 1988, pp. 5-14
14. Juang, J.-N., Applied System Identification, Prentice Hall, NJ 1994, pp. 175-182
15. Phan, M., Horta, L.G., Juang, J.-N., and Longman, R.W., *Improvement of Observer/Kalman Filter Identification (OKID) by Residual Whitening*, Journal of Vibration and Acoustics, Vol. 117, No. 2, 1995, pp. 232-239.