

Control of an RC Segway

Adam Cardi
Matt Wagner

Introduction

The aim of this project is to provide insight into the control theory behind the operation of a small remote controlled balancing robot. The robot's design was modeled after the Segway Human Transporter, in that it utilizes a mass imbalance actuated by a servo to control the vehicle's forward and reverse operation and utilizes a direct input to the controller for steering. The user can fully control the robot's motion with the use of a 2 channel R/C transmitter, and it can balance itself with no input from the user.



Figure 1: The completed model Segway

In order to control the robot, it is desired to have knowledge of as many states of the system as possible. In this particular case, the three states under consideration are angular position, angular rate and linear velocity. Angular position refers to the tilt of the model Segway's platform away from its neutral balance point, angular rate refers to the rate of change of this angle with respect to time, and velocity refers to the model Segway's linear velocity relative to the ground. Controlling the linear position is not of concern because any control gain applied to a position state would overwhelm the system since position can grow unbounded with normal driving.

In order to measure these states, an accelerometer is used to measure platform angle, and a rate gyro is used to measure the platform's angular rate. Once these signals are conditioned and amplified appropriately, they are fed to a Motorola HC11 microcontroller for computation. Efficient programming in assembly language resulted in a net sampling rate of 100 Hz, more than adequate for this project. The HC11 executes a stored control algorithm and outputs a pulse width modulated signal to two DC gearhead motors, one directly connected to each wheel. The HC11 has no knowledge of the mass imbalance of the system, just as the Segway has no direct knowledge of the rider leaning forward or back to control the system's speed. The HC11 does however have direct knowledge of the steering input, and must condition each motor's PWM signal appropriately to provide the desired output.

This report will cover the control theory used to make this model Segway system stable. Due to time limitations, only a simple control algorithm was implemented to the actual model. The theoretical aspects of this control law will be discussed and recorded data compared to theoretical plots. Deviations from the ideal model will be discussed, and other issues such as sensor dynamics and coupling will be explored. Furthermore, methods will be discussed to improve the system's performance. The application of two reduced order observers will be examined and their effects shown on the system's performance.

Model Assuming Ideal Measurements

The first step in the design process involved a simplified linear model with full state feedback. It was desired that the system keep its balance and maintain zero final velocity after being subjected to an impulsive rotational disturbance. A second type of disturbance arises when the model Segway is commanded to translate via the actuation of an eccentric mass. This type of disturbance should result in a finite steady state velocity, which will be used advantageously to drive forward and backward. A diagram of the system can be seen below in Figure 2.




Figure 2: The System diagram

State equations are developed from Newton-Euler equations of motion and are used to obtain matrices for the general form :

$$\dot{x} = Ax + Bu + Fd \quad (1)$$

Symbols Used

θ = platform deviation from vertical

R = wheel radius

L = distance from wheel centerline to center of mass

l = distance from wheel centerline to disturbance mass

d = horizontal displacement of disturbance mass

J = moment of inertia

α = motor constant 1

β = motor constant 2

D = duty cycle (range: $-1 \leq D \leq 1$)

M = system mass

m = disturbance mass

T_e = motor torque

k_t = torque constant

k_e = back emf constant

r_a = armature resistance

V_s = motor voltage

τ = time constant of lowpass filters

C_1 = Rotational Damping

C_2 = Linear Damping

Equations of Motion

$$\sum M_o = J\ddot{\theta} = MgL\theta + mgl\theta - T_e - C_1\dot{\theta} \quad (2)$$

$$\sum F_x = (M + m)\ddot{x} = \frac{T_e}{R} - C_2\dot{x} \quad (3)$$

where:

$$\sin(\theta) \approx \theta$$

$$T_e = \alpha D - \beta \dot{x}$$

$$\alpha = \frac{k_t V_s}{r_a}$$

$$\beta = \frac{k_t k_e}{R r_a}$$

Simplifying these equations and rewriting in state variable form yields:

$$\ddot{\theta} = \frac{(MgL + mgl)}{J} \theta + \frac{\beta}{J} \dot{x} - \left(\frac{\alpha}{J}\right) D + \frac{mg}{J} d - \frac{C_1}{J} \dot{\theta} \quad (4)$$

$$\ddot{x} = \frac{\alpha}{R(M+m)} D - \frac{(\beta + C_2)}{R(M+m)} \dot{x} \quad (5)$$

In matrix form these equations are represented by:

$$\frac{d}{dt} \begin{pmatrix} \theta \\ \dot{\theta} \\ \dot{x} \end{pmatrix} = \begin{pmatrix} 0 & 1 & 0 \\ \frac{MgL + mgl}{J} & -\frac{C_1}{J} & \frac{\beta}{J} \\ 0 & 0 & \frac{-(\beta + C_2)}{R(M+m)} \end{pmatrix} \begin{pmatrix} \theta \\ \dot{\theta} \\ \dot{x} \end{pmatrix} + \begin{pmatrix} 0 \\ \frac{-\alpha}{J} \\ \frac{\alpha}{R(M+m)} \end{pmatrix} D + \begin{pmatrix} 0 \\ \frac{mg}{J} \\ 0 \end{pmatrix} d \quad (6)$$

While most parameters were easily obtained, some required more effort. The system's moment of inertia was not immediately obvious, and was measured indirectly. The wheels of the platform were removed and the model Segway hung upside down from its axles. The system's period of oscillation was timed and used to back-calculate the system's moment of inertia. This was verified by changing the sign of the (3, 2) entry in the A matrix in Equation (6), which made the open loop system oscillate in a stable fashion. A graph of this can be seen below in Figure 3. (This would be analogous to the having the system's center of mass below the axis of rotation) It should be noted the model closely matches the system's actual period of oscillation of 1.14 s.




Figure 3: Verification of moment of inertia.

Other parameters that required investigation were the motor characteristics. The back EMF constant was experimentally determined and verified with data from the manufacturer's website. It was determined by spinning one motor at a known speed (60 RPM) and measuring the voltage across the terminals. The back EMF constant was found to match the manufacturer's published rating closely. The torque constant was assumed to equal the back EMF constant per manufacturer's ratings, a common assumption for permanent magnet DC brush motors this size.

Simulation of the Analytical Model

Before discussing the model, a few comments are needed about additions to the actual model Segway necessary to mitigate problems encountered in implementation. The first attempt at controlling the system proved unsuccessful, primarily due to physical feedback to the sensors. The sensors were initially fed through a 100Hz RC lowpass filter to remove any high frequency noise from the signal. This filtering performed as expected and output was verified on an oscilloscope. The system also produced a smooth output to the motors while held in midair, but immediately developed uncontrollable oscillation when placed on the ground. It was discovered that self-excitation was occurring since the sensors were physically coupled to the platform in the translational direction even though they are only meant to measure rotational quantities. The accelerometer was particularly sensitive to platform vibrations induced by the motors. These oscillations were found to be at a frequency of approximately 50Hz, and it was determined that more filtering was necessary.

The filtering cutoff frequency was lowered until the problem was alleviated, with a final cutoff frequency of 5Hz determined to be optimum. Faster cutoff frequencies reintroduced oscillations, and slower cutoff frequencies yielded measurements with an unacceptable phase delay.

For the purposes of data comparison, these sensor dynamics were modeled as they introduce a significant phase lag into the feedback loop. The transfer function for a lowpass RC filter equal to:

$$H(s) = \frac{1}{\tau s + 1}, \text{ where } \tau = RC \quad (7)$$

Another issue that merits discussion is the placement of the center of mass of the system. The largest components of the system's mass were the batteries, so there was freedom to place the center of mass where desired. Initial intuition pointed to placing them as far from the rotation axis as possible because an object with a high moment of inertia seemed easier to balance, but this was proved false after much trial and error. The batteries were eventually moved to a location as close to the axles as possible, which made the system immensely easier to work with. At this point the importance of having a linear plant in a linear controls system was realized, a point that was overlooked with the initial assumption to place the mass as high as possible. The $\sin(\theta)$ term used in the moment

equation is very nonlinear when the height of the center of mass is high, and the effect of this nonlinear term made the system difficult if not impossible to control adequately using linear techniques.

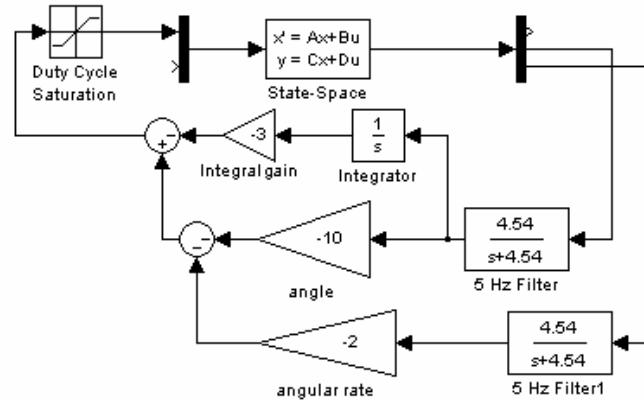


Figure 4; The block diagram for PD and PID control.

The first attempt at controlling the system was accomplished by using direct feedback from the angular position and angular rate sensors, effectively providing PD control. A block diagram of this system can be seen in Figure 4. The PD system shares the same block diagram as the PID system, with the integral gain in the PD case set to zero. In the theoretical model, pole placement will yield appropriate gains based on the desired response of the system. In reality it was realized that with the system's limitations, only a small range of gains was possible. Gains too large still introduced instability to the system, and the processor's digital nature limited gain selection to a small usable range. The system was stable with PD control, however the response to a disturbance was undesirable as predicted with the analytical model.

A few observations can be made about the system's performance with PD control. The model predicts the system will remain upright when subject to an impulsive disturbance but develop a velocity error which will increase with time. This was precisely the response observed, and the reason why methods beyond PD control were explored. The sensor outputs were read by an oscilloscope and used to capture the system's response while subjected to a disturbance. Plots of this data and output of the analytical model can be seen in Figure 5.




Figure 5: The actual and theoretical Data for PD control.

The model predicts an initial settling time after the disturbance, with the steady state response exhibiting a small static angle and a velocity error with a constant slope. The velocity error's constant slope will yield a constant acceleration, which counteracts the moment placed on the system by steady state angle error. The actual system exhibits a similar response, with the exception of its continued oscillation. This is most likely due to the fact that a steadily increasing velocity is difficult to obtain in practice, as any slight disturbance to the system (irregular surface, friction) will result in a disturbed angle and thus reintroduced oscillation. The idea of strict PD control was soon abandoned, and methods to account for steady state error in angle explored.

Since full state feedback including velocity was not possible with available hardware, integrating the angle to provide PID control was implemented in software. This integral was added to the analytical model, and as expected drove the angle error and thus velocity error to zero.

In reality, the model had a difficult time achieving an integral of zero, often resulting in a steady oscillation around the balance point. The digital nature of the system made it difficult to drive the integral to zero, especially at small angles of rotation. Thus, a threshold was implemented close to the balance point within which the integral would not grow but would be decremented. This allows an excess integral to be decreased if the system is near the balance point, and prevents an unnecessary integral from accumulating if one is not needed. Despite these implementation differences, the system matches the predicted model quite well.

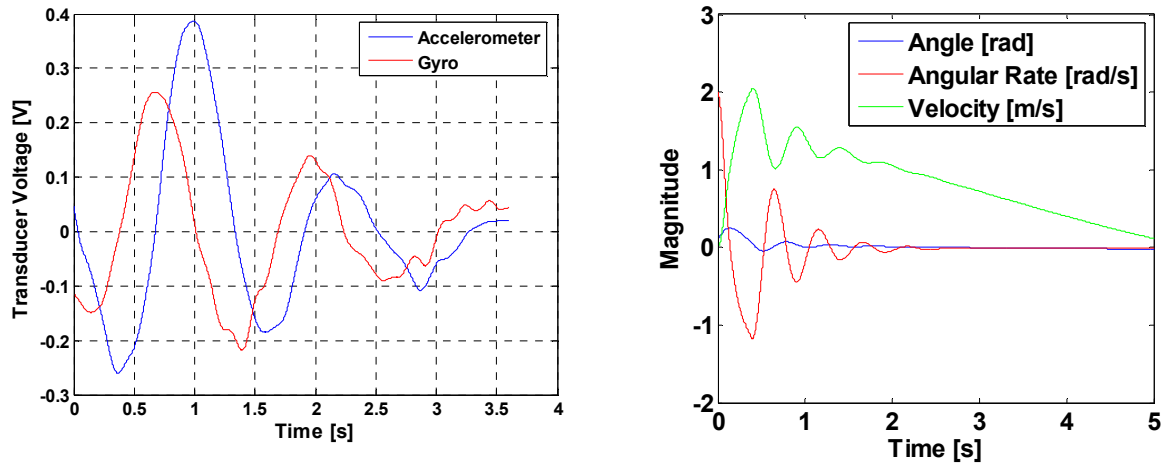


Figure 6: The actual and theoretical Data for PD control.

Comparing the analytical graph to that of PD control shows one immediately obvious difference. Where PD control resulted in a steadily increasing velocity, PID control shows velocity approaching zero. The system behaved as such, and responded well to a disturbance with angle and velocity quickly approaching zero. It should be noted that the analytical model's velocity will continue to oscillate for time after 5s, until the integral is eventually driven to zero. The actual model settles much faster, largely due to the integral being forced to zero within the threshold in software. The analytical model shows the system oscillating initially while responding to a disturbance, followed by a period of no oscillation while velocity slows to zero. This response can be observed in the real system, both in the recorded data and in video documentation of a disturbance input. It can also be noted that the response of this system more closely matches the output of the analytical model, largely due to the system's reduced velocity. Since velocity approaches zero more quickly, the opportunity for a surface irregularity or other disturbance to upset the system is reduced, thus resulting in fewer unpredictable oscillations.

The need for more sophisticated control

At this point, the model exhibited acceptable performance, could easily balance itself and was very controllable by a user via radio control. Due to time restrictions, these were the only control algorithms implemented to the actual system. However, it was desired to have more control over the system's response.

With only information about angular position and angular rate available, there was not as much flexibility to shape the system's dynamic performance. In practice, the controller gains needed to be high enough to make the system stable, but not too high as to make the system oscillate uncontrollably. This window was quite small, thus providing little control over the system's characteristics. It was observed in the analytical model, and expected from knowledge of state space techniques, that if velocity were known (thus providing full state feedback) control of the system's response could be greater. Aside from physically measuring velocity, an observer can be constructed to estimate it from

the directly measured states of angular position and angular rate. This technique was explored as an academic exercise to try and improve the system's performance without requiring any further sensors.

Full State Feedback with a Reduced Order Observer Assuming Ideal Measurements

A reduced order observer is to be designed to estimate x_2 (in this case, solely the velocity state) with the following form:

$$\hat{x}_2 = Ly + z \quad (8)$$

$$\dot{z} = Fz + \bar{G}y + Hu \quad (9)$$

$$H = B_2 - LC_1B_1 \quad (10)$$

$$F = A_{22} - LC_1A_{12} \quad (11)$$

$$\bar{G} = (A_{21} - LC_1A_{11})C_1^{-1} \quad (12)$$

The dynamics of the observer are determined by the eigenvalues of the F matrix. It was desired that the observer poles be 5-10 times as fast as the fastest pole of the plant, and this was accomplished by choosing appropriate gains in the L matrix. With the A matrix partitioned such that the directly measured states (angular position and rate) occupy the first two rows, the following can be shown to represent the observer gains:

$$H = \frac{\alpha}{R(M+m)} + \frac{l_2\alpha}{J} \quad (13)$$

$$G = \left[\begin{array}{cc} \frac{l_2(MgL+mgl)}{J} & \frac{l_2C_1}{J} - l_1 \end{array} \right] \quad (14)$$

$$F = \frac{-\beta - C_2}{R(M+m)} - \frac{l_2\beta}{J} \quad (15)$$

where:

$$L = [l_1 \quad l_2] \quad \text{and} \quad C1 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$$

Since the eigenvalue of F equals the value of F, it can be noted that only the selection of l_2 will affect the dynamics of the observer.

The structure of the observer in block diagram form, as it was implemented in Simulink is shown below in Figure 7.

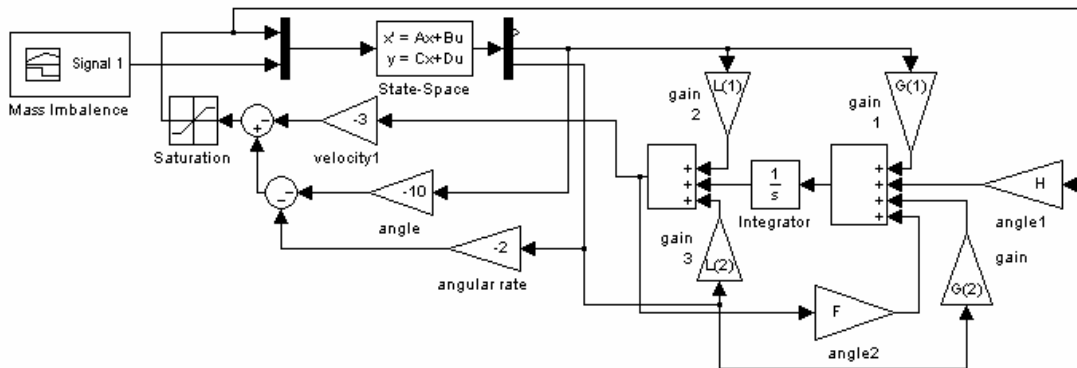


Figure 7: Full state feedback with a reduced order observer estimating one state.

For this simulation, the measurements made on the state of angular velocity and angular position were assumed to be ideal, and they were used in the reduced order observer to estimate the linear velocity of the cart. It should be noted that there is saturation block on the control signal going into the system. This is because the input to the system is a duty cycle, bounded between the values of -1 and 1. Because the dominant poles of the plant have natural frequencies of about 1 Hz the single pole of the observer was chosen so that it was at least 5 times faster. The eigenvalue of the first order observer is -26.3.

As discussed previously, it was determined empirically on the model Segway that there was a relatively small region of acceptable gains on angular position and angular rate. Therefore in choosing a gain for velocity, it was decided that the gains of angular position and angular rate would be held constant, and system performance as a function of varying gain on angular velocity would be investigated. Below in Figure 8 is a root locus for a varying gain on the linear velocity state:

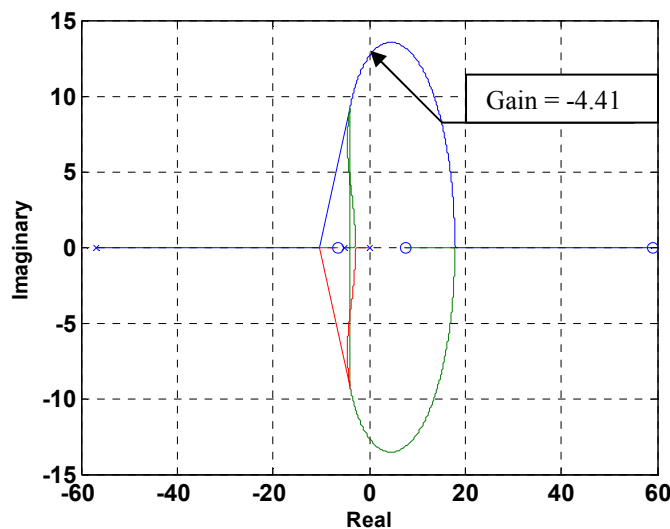


Figure 8: The root locus plot for the system with ideal measurements as a function of velocity gain.

The range of gains on linear velocity that provided closed loop stability was $0 > \text{gain} > -4.41$. It was determined that a gain of -3 provided the right balance of fast response time and damping for the dominant pair of second order poles in the system.

Below in Figure 9 are the responses and controller input duty cycle when the system is subjected to an initial condition of 115 degrees per second on the angular velocity of model Segway; this is meant to simulate the response to an impulsive type disturbance.




Figure 9: Response to an initial condition on angular velocity.

It should be noted that all three states approach zero with minimal oscillation. The nonlinear effect of the duty cycle saturating in the initial response does not significantly affect the dynamics of the system because with the saturation of the duty cycle removed the responses look quite similar. This type of response is desirable both in the case of the radio controlled model Segway as well as the commercially available Segway. When impulsive type disturbances are encountered, such as driving over bumps in the road, it is important that the angle and angular rate are stabilized quickly. If this observer was implemented on the actual remote control Segway, it is anticipated that the actual responses would be different from those above because there are no filter dynamics in this model. Adding the filter dynamics to the Simulink model shows the observer performs unacceptably, since the observer has no knowledge that the signals it uses are not the true values from the system. In actuality, first order low pass filters with cutoff frequencies of 4.5 Hz were used on the measured states of angular rate and angular position in order to prevent the plant from exciting itself, as well as to mitigate the effects of electromagnetic interference. For this reason, this simplistic approach would likely not work unless the sensors could be filtered with a higher cutoff frequency.

Full State Feedback with a Reduced Order Observer Assuming Filtered Measurements

An observer which takes into account the filter dynamics of the system can be implemented using the same technique as the observer estimating one measurement. This observer knows the data used for its input has been altered by each filter, and thus can calculate what the actual unfiltered state of the system is. The system can then use this data for full state feedback, as though it was essentially unaffected by the filter's detrimental phase lag.

The A matrix must be expanded as shown, and partitioned appropriately. With the filters defined by the transfer function in Equation (7) the state space model for the system becomes:

$$\frac{d}{dt} \begin{pmatrix} \theta_f \\ \dot{\theta}_f \\ \theta \\ \dot{\theta} \\ \dot{x} \end{pmatrix} = \begin{pmatrix} -\frac{1}{\tau} & 0 & \frac{1}{\tau} & 0 & 0 \\ 0 & -\frac{1}{\tau} & 0 & \frac{1}{\tau} & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & \frac{MgL+mgl}{J} & -\frac{C_1}{J} & \frac{\beta}{J} \\ 0 & 0 & 0 & 0 & \frac{-(\beta+C_2)}{R(M+m)} \end{pmatrix} \begin{pmatrix} \theta_f \\ \dot{\theta}_f \\ \theta \\ \dot{\theta} \\ \dot{x} \end{pmatrix} + \begin{pmatrix} 0 \\ 0 \\ 0 \\ -\frac{\alpha}{J} \\ \alpha \\ \frac{\alpha}{R(M+m)} \end{pmatrix} D + \begin{pmatrix} 0 \\ 0 \\ 0 \\ \frac{mg}{J} \\ 0 \end{pmatrix} d \quad (16)$$

where θ_f and $\dot{\theta}_f$ are the filtered versions of θ and $\dot{\theta}$ respectively, and τ is the time constant of the first order low pass filter.

This system is a far more accurate representation of the true dynamics of the model Segway because it was observed empirically that the behavior of the system was intimately connected with the cutoff frequency of the filters. With this new model the measured states will be θ_f and $\dot{\theta}_f$, just as they were in reality, and they will be used to estimate θ , $\dot{\theta}$, and \dot{x} .

Computing the F, G and H matrices is an exercise in matrix algebra, with the only freedom for design being present in the selection of gains in L.

$$L = \begin{bmatrix} l_1 & 0 \\ 0 & l_2 \\ l_3 & l_4 \end{bmatrix} \quad (17)$$

The selection of gains in L shows two terms equal to zero, which is implemented in an attempt to reduce the number of degrees of freedom involved in selecting L values. It was desired that θ_f have influence on the estimate of θ , and that $\dot{\theta}_f$ have influence on

$\dot{\theta}$. Also, it was desired that the velocity estimate be influenced by both θ_f and $\dot{\theta}_f$, as indicated by the last row of the L matrix. The C_1 matrix remains as in the single estimate case.

Solving for F, G and H yields:

$$H = \begin{bmatrix} 0 \\ -\frac{\alpha}{J} \\ \alpha \\ \frac{\alpha}{R(M+m)} \end{bmatrix} \quad (18)$$

$$G = \begin{bmatrix} \frac{l_1}{\tau} & 0 \\ 0 & \frac{l_2}{\tau} \\ \frac{l_3}{\tau} & \frac{l_4}{\tau} \end{bmatrix} \quad (19)$$

$$F = \begin{bmatrix} -\frac{l_1}{\tau} & 1 & 0 \\ \frac{(MgL+mgl)}{J} & -\frac{C_1}{J} - \frac{l_2}{\tau} & \frac{\beta}{J} \\ -\frac{l_3}{\tau} & -\frac{l_4}{\tau} & \frac{(-\beta-C_2)}{R(M+m)} \end{bmatrix} \quad (20)$$

Placing the eigenvalues of F is not as simple as in the last case, and solving symbolically proved a formidable task even for a computer. The values of the constants are known, and numerically solving for the eigenvalues of F yields the characteristic equation:

$$s^3 + (8.18 + 4.54l_1 + 4.54l_2)s^2 + (-10.62 + 37.21l_1 + 26.81l_2 + 20.66l_1l_2 + 46.22l_4)s + (-142.29 + 61.37l_1 + 121.87l_1l_2 + 46.21l_3 + 210.1l_1l_4) = 0 \quad (21)$$

Again, it was desired that the observer poles were 5-10 times as fast as the fastest pole of the plant. Through trial and error, the gains were modified while monitoring the eigenvalues, with the final L matrix as follows:

$$L = \begin{bmatrix} 20 & 0 \\ 0 & 10 \\ 10 & 10 \end{bmatrix} \quad (22)$$

The structure of this system is shown below in Figure 10.

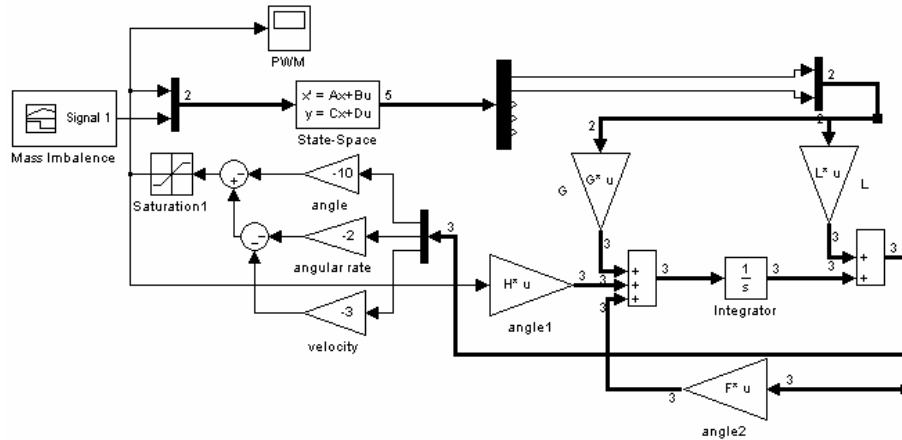


Figure 10: The system with reduced order observer estimating three states.

The gains of the L matrix in Equation (22) yield observer poles at: -91.6 , and $-26.3 \pm 13.5i$. The real part of the slowest pole in this observer is the same as the single pole in the previous observer where one state (velocity) was being estimated from ideal measurements; this is done for comparison purposes. In the previous case where no filters were present, the gains on the estimated velocity were determined from a root locus analysis. Below in Figure 11 is the root locus plot for a varying gain on the linear velocity state for the system with filter dynamics given by Equation (16).

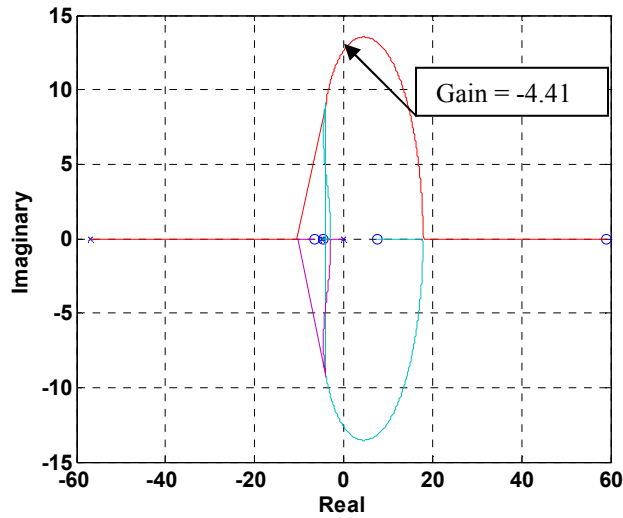


Figure 11: The root locus plot for the system with ideal measurements as a function of velocity gain.

It should come as no surprise that the filters have no affect on the range of stability of the system, however the presence of the extra zero will surely affect the transient response. The range of stability on a gain proportional to velocity is still $0 > \text{gain} > -4.41$, Once again, a gain of -3 on the velocity state proved to be the right balance between damping and response time of the dominant poles in the system. To compare the effects of the new filters, this new system with filtered measurements was subjected to the same initial condition on angular velocity as the previous system, which had ideal measurements.




Figure 12: Response to an initial condition on angular velocity.

It should be noted that in this case the response is slower and all states have larger amplitude oscillations than in the case where there were no filters on the measurements. This can be attributed to the fact that the low pass filters prevent high frequency content from getting to the observer so therefore the predicted states are not as accurate.

Next, the system was subjected to a disturbance which would simulate a rider leaning forward (in order to drive the vehicle forward), or in the case of the model Segway, the actuation of a mass imbalance by a servo controlled through a radio control link. This input appears in Equations (16) and (6) as d , the distance to the moment arm of the mass imbalance. In Figure 13 is the input d to the system as a function of time:




Figure 13: The input to the system that simulates driving forward.

This input signal can be interpreted as a command to go forward at $t = 1$ s and then a command to stop at $t = 3.5$ s. It was decided to use “ramps” rather than “steps” because in the model Segway, or in the case of the real Segway, the mass is shifted gradually. The response of the model is shown below in Figure 14 when subject to the input shown in Figure 13.




Figure 14: The system response when it is commanded to drive forward.

It should be noted that when the model is commanded to drive forward by the mass imbalance at $t = 1$ s, the angle settles in about 2 s to about 3.8 degrees, which allows for a constant velocity of 0.12 m/s. This response would be acceptable if it were implemented on the model Segway.

Conclusion

Clearly, the three state estimate observer adds a new element of control and tunability to the model Segway. It should be noted that the observer does an excellent job of estimating the actual states of the system despite having its input subject to heavy filtering and significant phase lag at higher frequencies. If implemented, the three state observer has the potential to make the model Segway a robust and smooth system.

The three estimate observer, although simple in block diagram form, would be complex to implement in hardware. Many elements are matrices, and for simplicity the most practical way to implement this type of observer would be in software. This project's use of a Motorola HC11 would limit this approach, but many newer microprocessors would perform acceptably.

If the physical coupling problem of the sensors could be alleviated, the single estimate observer could potentially be used with similar results. The single estimate observer did not react well in simulation to integration with the current RC filters, which unfortunately at this time are elements that cannot be removed. However, in the event that the sensors could be isolated sufficiently from the platform to reduce the need for heavy filtering, the single estimate observer could be a simple approach to achieve significant performance gains.

Since the single estimate observer only has 9 elements, it would be an interesting academic exercise to build an analog computer to do its job. Summers, integrators and gains can all be easily implemented using opamps. Careful attention would need to be given to all parts of the circuit to avoid saturation, but in practice this approach would very likely be feasible given enough time.

This project was a beneficial learning experience in a number of ways. The single most important lesson learned was to improve the plant as much as possible rather than account for its shortcomings in software. Also, real world issues of noise and sensor coupling were experienced firsthand. These issues can change a textbook system into one that requires more creative effort to control. Lastly, the notion that computed gains from a model can be simply applied to a real system was realized to be false. Regardless of the depth and complexity of a given model, the real system will most likely need to be tuned to some degree before completion. In this respect, the practical aspect of controls can be viewed as an art form as much as an exercise in pure mathematics.