

The Atlantis Project: A GPS-Guided Wing-Sailed Autonomous Catamaran

GABRIEL HUGH ELKAIM

Autonomous Systems Lab, Computer Engineering, University of California, Santa Cruz,
Santa Cruz, CA 95064 USA

Received September 2006; Revised January 2007

ABSTRACT: *An autonomous catamaran, based on a modified Prindle-19 day-sailing Catamaran, was built to test the viability of GPS-based system identification for precision control. The catamaran was fitted with several sensors and actuators to characterize the dynamics. Using an electric trolling motor, and lead ballast to match all-up weight, several system identification passes were performed to excite system modes and model the dynamic response. LQG controllers were designed based on the results of the system identification passes, and tested with the electric trolling motor. Line following performance was excellent, with cross track error standard deviations of less than 0.15 meters. The wing-sail propulsion system was fitted, and the controllers tested with the wing providing all forward thrust. Line following performance and disturbance rejection were excellent, with the cross track error standard deviations of approximately 0.30 meters, in spite of wind speed variations of over 50% of nominal value.*

INTRODUCTION

This paper details the progress of the Atlantis project, pictured in Figure 1, which began with the conception of an unmanned, autonomous, GPS-guided, wing-sail-propelled sailboat. The Atlantis project has been very much a “systems” approach, with substantial innovations in the areas of wind-propulsion, overall system architecture, sensors, system identification and control.

Functionally, the Atlantis is the marine equivalent of an unmanned aerial vehicle, and would serve similar purposes. The Atlantis project has been able to demonstrate an advance in control precision of a wind-propelled marine vehicle from typical commercial autopilot accuracy of 100 meters to an accuracy of better than one meter. This quantitative improvement enables new applications, including unmanned station-keeping for navigation or communication purposes, autonomous “dock-to-dock” capabilities, emergency “return unmanned” functions, and many others still to be developed.

The wind-propulsion system is a rigid wing-sail mounted vertically on bearings to allow free rotation in azimuth about a stub-mast. Aerodynamic torque about the stub-mast is trimmed using a flying tail mounted on booms joined to the wing. This arrangement allows the wing-sail to automatically

attain the optimum angle to the wind, and weather vane into gusts without inducing large heeling moments. Modern airfoil design allows for an increased lift-drag (L/D) ratio over a conventional sail, thus providing increased thrust while reducing the overturning moment.

The system architecture is based on distributed sensing and actuation, with a high-speed digital serial bus connecting the various modules together. Sensors are sampled at 100 Hz, and a central guidance navigation and control (GNC) computer performs the estimation and control tasks at 5 Hz. This bandwidth has been demonstrated to be capable of precise control of the catamaran. The distributed architecture is both more robust and less expensive than systems that employ a high-speed, and often analog, star-configuration topology with centralized sensor interpretation and actuation.

The sensor system uses differential GPS (DGPS) for position and velocity measurements, augmented by a low-cost attitude system based on accelerometer- and magnetometer-triads. Accurate attitude determination is required to create a synthetic position sensor that is located at the center-of-gravity (CG) of the boat, rather than at the GPS antenna location.

Experimental trials recorded sensor and actuator data intended to excite all system modes. A system model was assembled using Observer/Kalman System Identification (OKID) techniques. An LQG controller was designed using the OKID model,




Fig. 1—The Atlantis, based on a Prindle-19 Catamaran, with a self-trimming wingsail. The Atlantis was designed to demonstrate a very high precision of navigation and control, even in the presence of wind and waves. As shown here, testing in Redwood City harbor, January 2001, the Atlantis was able to achieve line following to better than 30 cm. 1- σ under wind propulsion. The author and colleagues on board are human ballast to prevent capsizing during this initial test.

using an estimator based on the observed noise statistics. Experimental tests were run to sail on a precise track through the water, in the presence of currents, wind and waves.

SYSTEM DESCRIPTION

In order to experimentally validate the concepts presented in this research, a prototype system was built based on a heavily modified Prindle-19, day-sailing catamaran. The catamaran was 7.2 m long, 3 m wide, and was originally equipped with a sloop rig sail with 17 m² of sail area. Directional control is based on rudders at the end of each hull, and retractable centerboards approximately $\frac{1}{2}$ m behind the main crossbeam. Several sensors and actuators were installed within the hulls, and the entire sailing system (mast, boom, main and jib sails) was replaced with a vertical self-trimming wing (wing-sail) suspended on spherical roller bearings.

There are several main subsystems on the Atlantis, and all of them are connected to each other via a high-speed serial network. The network utilized is the Controller Area Network (CAN) bus, which was designed by Bosch electronics for robust component communication in an automotive environment [1]. The entire wiring bus on the Atlantis consists of four (4) wires: power (+12V), ground, CAN_hi, and CAN_low. There are many advantages to this setup, but the ease of troubleshooting and flexibility of physical topology are at the forefront of utility in this design.

Figure 2 shows the system as it ties together logically and electrically on the CAN bus. The main subsystems are: attitude system, anemometer, hullspeed, rudder angle, rudder actuator, GPS

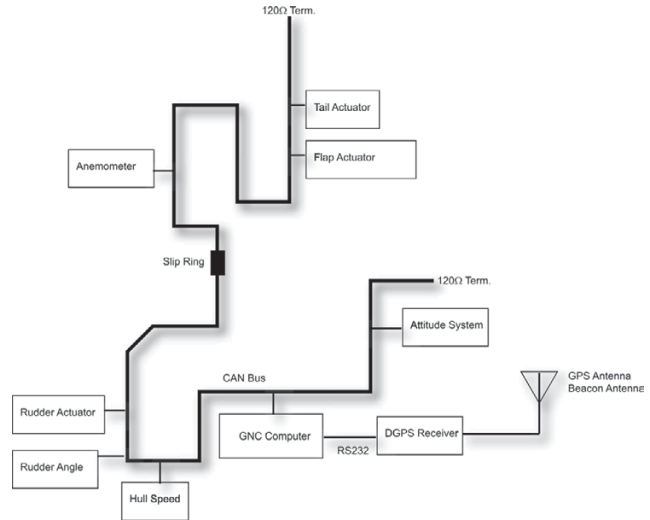


Fig. 2—Block diagram of the major subsystems of the Atlantis. Note that both power and CAN signals go through a slip ring at the top of the stub mast in order to electrically connect the rotating wing to the rest of the system.

receiver, and wing-sail. These subsystems communicate to the main GNC computer that computes the current estimate of the state, and returns the required commands to the actuator in order to achieve control.

The attitude system, pictured in Figures 3 and 4, consists of a three-axis magnetometer, a two-axis accelerometer, and a Siemens 515 microcontroller. It functions based on a novel gyro-free quaternion based solution to the vector matching problem first proposed by Wahba in 1966 [2]. The algorithm is discussed extensively in [3, 4, 5]. The attitude system is mounted alongside the Global Positioning System (GPS) antenna, inside a waterproof Pelican box on a wooden crossbeam at the forward stay location (note that the wooden crossbeam was added for increased structural rigidity of the hulls, due to the stresses induced by the wing).

The GNC computer, shown in Figure 5, a Pentium class laptop, is placed in another waterproof case,




Fig. 3—The Atlantis uses a low-cost attitude solution that is based on a quaternion solution to Wahba's problem, where the two observed vectors are Earth's gravitation and magnetic fields. The solution produces 100 Hz. attitude data and is based on an Analog Devices 2-axis accelerometer and Honeywell 3-axis magnetometer.




Fig. 4—A close-up of the low-cost attitude solution used on the Atlantis, and pictured in Figure 3. The ruler is there for scale, and shows that the entire board is less than 6 inches long and 4 inches across. The microcontroller is an Infineon 515 processor, and the attitude is being relayed on the CAN bus.



Fig. 5—Guidance, Navigation, and Control computer box, with the Trimble Ag122 DGPS Receiver, A Pentium Class laptop, and DC-DC converter to power the laptop, and an ESD Parallel port CAN dongle. The entire box is sealed and has water-tight connections for Power, CAN, and the RF signal from the GPS antenna.

along with the Trimble Ag122 GPS receiver. The GNC computer is equipped with an ESD parallel port dongle that allows communication over the CAN bus. A DC/DC converter insures that the laptop draws power from the boat power bus rather than its own internal batteries.

Inside the starboard hull, beneath the rear inspection cover, are two Infineon 505 microcontrollers, one for the hullspeed and rudder angle sensor, and the other for the rudder actuator, pictured in Figure 6. There is a Standard Communications Electronics marine through-hull speed sensor, pictured in Figure 7, in the bottom of the starboard hull, and a LoHet magnetic field effect sensor between two magnets on the rudder hinge line to measure rudder angle. The actuator is a fractional horsepower DC motor, with a lead screw assembly, constrained to rotate only in yaw, and an Infineon H-bridge mosfet drive for electronic control of the motor.

Inside the stub-mast, a Mercotac slip ring allows for a full 360 degree rotation without twisting the four wires that comprise the wiring harness of the Atlantis. The wing itself is built in three sections that are assembled on site. The lower section contains an electronics pod with the batteries, ballast, and battery-charging electronics. A microcontroller and DC motor are used to control the trailing edge flap. It also contains the anemometer microcontroller, with the Standard Communications Electronics marine transducer head, pictured in Figure 8, attached to the top of the electronics pod lid.

Within the wing are four actuators that are identical to the rudder actuator, which actuate the




Fig. 6—Rudder actuator is made from a Pittman 24V DC Motor turning a lead-screw. The motor has a 512 line per revolution quadrature phase encoder attached, and is directly attached to the lead screw. With this arrangement, a slew-rate of over 20 degrees per second was achieved.



Fig. 7—Hullspeed sensor is located towards the rear of the starboard hull, and is made by the Standard Communications Electronics Corporation. Due to its location behind the centerboard of the hull, the signal is quite noisy from the turbulence.



Fig. 8—A 3-cup anemometer is used to determine wind speed and direction. The unit is made by Standard Communications Electronics, and has a two-pulse per revolution output for the cups, and a quadrature analog signal for the direction of the weathervane. It was calibrated in a wind tunnel.

trailing edge flaps and the tail. The moment balance between the wing and the tail keeps the wing-sail at a constant angle of attack relative to the wind. As long as the wind does not cross the centerline of the boat, then the wing continues to provide thrust in the correct direction for forward motion through passive stability of the wing-sail system. Should the wind cross through the centerline of the boat, then the position of the flap and tails must be reversed, which corresponds to tacking or jibing depending on whether the wind crosses the centerline facing aft or forward, respectively.

The anemometer is used to measure the wind speed and direction relative to the angle of the wing-sail. Essentially, this is a measure of angle of attack of the wing.

SYSTEM IDENTIFICATION METHODOLOGY

In order to control the Atlantis, a system model needed to be assembled. While several good modeling techniques exist to model a powered boat through the water [6], they remain complicated and difficult to calculate. In order to reduce the model order, and obtain a model that would have sufficient fidelity for active control, several different methodologies were attempted. In order to formulate the equations of motion, the Atlantis is assumed to be traveling upon a straight line, conveniently assumed to be coincident with the X-axis, through the water at a constant velocity, V_x . The distance along that line is referred to as X , the along track distance. The perpendicular distance to the line is referred to as Y , the cross track error, and the angle that the centerline of the Atlantis makes with respect to the desired path is defined as Ψ , the angular error. Figure 9 illustrates the mathematical model of the assumed path of the Atlantis.

This first pass model is a simple kinematic model that assumes that the rudders cannot move sideways through the water. This places a kinematic constraint upon the motion of the entire boat, and the linearized analysis produces the following continuous time state-space equations:

$$\begin{bmatrix} \dot{Y} \\ \dot{\psi} \\ \dot{\delta} \end{bmatrix} = \begin{bmatrix} 0 & V_x & 0 \\ 0 & 0 & \frac{V_x}{L} \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} Y \\ \psi \\ \delta \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} u \quad (1)$$




Fig. 9—The basic equations of motion for the Atlantis on the water. A simplified model is used as a basis for understanding the vehicle motion. Velocity is along the X-axis, with cross track error being measured in the Y-axis. The heading error is measured from the X-axis to the centerline, and defined as Ψ .

Where Y , ψ , and δ are defined as the cross track error, azimuth error, and rudder angle, and L is the length from the C.G. of the boat to the center of pressure of the rudders. Note that these simplified equations of motion are insufficient to control the boat to great precision, but are excellent for generating intuition for the system identification process. Equation 1, when cast into transfer function form, becomes a triple integrator, and cannot be stabilized by simple proportional control. In addition, the assumption of constant V_x is poor, since unless the wind can be controlled, the velocity will always be dependent on the speed of the wind. Closer inspection of Equation 1 shows that the errors in azimuth and cross track integrate not with time, but rather with distance traveled forward. What this means is that if the boat is sitting still in the water, no amount of rudder deflection will cause the azimuth to change. Likewise, when moving very quickly through the water, only very small inputs are required to turn the boat through a considerable angle.

Thus, recharacterizing the variables of interest, to make the system velocity invariant, both the azimuth and cross track error are normalized by velocity. The new variables of interest become \tilde{y} and $\tilde{\psi}$, with δ remaining the same as previously defined. The new simplified equations of motion become:

$$\begin{bmatrix} \dot{\tilde{Y}} \\ \dot{\tilde{\psi}} \\ \dot{\delta} \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & \frac{1}{L} \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \tilde{Y} \\ \tilde{\psi} \\ \delta \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} u \quad (2)$$

where

$$\begin{aligned} \tilde{Y} &= \frac{Y}{V_x} \\ \tilde{\psi} &= \frac{\Psi}{V_x} \end{aligned} \quad (3)$$

Note that this causes the state transition matrix to be constant. This concept of velocity invariance has been tested extensively on GPS-controlled farm tractors [7], and shown to work very well. Functionally, this means that controller design is reduced to a single (non-gain-scheduled) controller that automatically adjusts for the changing velocity based on decreasing input ranges as velocity increases.

In order to gather data to perform a proper system identification of the Atlantis, a series of open-loop line-following tests were conducted in which a human driver, through the Guidance, Navigation, and Control (GNC) computer, caused the rudders to either slew left or right at the maximum slew rate (25 degrees/s). Also, the driver commanded

the rudder slew rate to zero through the rudder actuator in order to track a roughly straight line. This “pseudo”-random input was designed to apply the maximum power to the Atlantis 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 in Figure 10.

Once the System Identification process is completed (and a state space representation of the system has been generated), a controller was designed using a standard Linear Quadratic Regulator (LQR) methodology. The quadratic cost, as calculated below in Equation 4, minimizes the weighted sum of the outputs (y_{max} and u_{max} are design parameters).

$$J = \sum_{k=0}^{\infty} \left(\tilde{x}_k^T C^T \begin{bmatrix} \frac{1}{y_{max}^2} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix} C \tilde{x}_k + \tilde{u}_k^T \begin{bmatrix} 1 \\ u_{max}^2 \end{bmatrix} \tilde{u}_k \right) \quad (4)$$

Note that the quantity $\tilde{x}_k^T C^T$ is nothing more than y^T which is $[y \ \psi \ \delta]^T$. In other words, the cost that is being minimized by the LQR controller is the weighted sum of the square of the crosstrack error and the input (rudder rate). This controller places no penalty on either heading error or rudder angle; it only tries to reduce crosstrack error.

In terms of system identification, this becomes useful as the inputs are scaled by velocity before being assembled into the system identification algorithm. Thus, the identified system is one that is the best model for the velocity invariant control.

The system identification methodology used for this work is the **Observer Kalman IDentification**

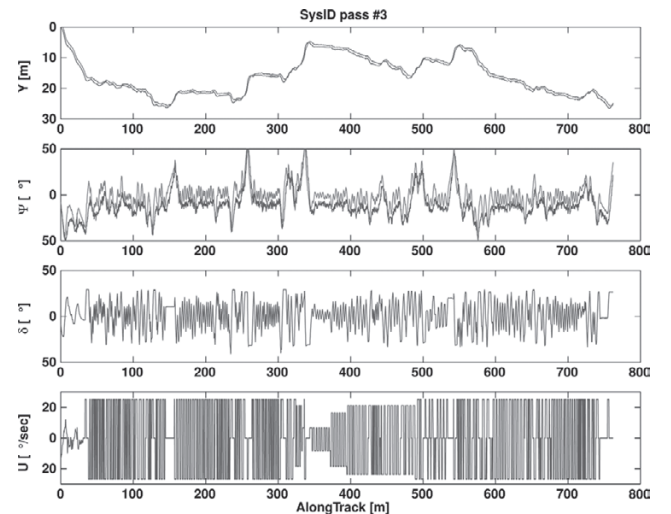


Fig. 10—A typical system identification pass. The input rudder angle slew rate, u , was input by a human driver attempting to keep the Atlantis on a roughly straight line. Note that this pass is over 700 meters long.

(OKID) [8]. Among its many advantages is a formulation that presumes a discrete-time, linear 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 [9]. The following discussion is a summary of the OKID algorithm, and is presented here for completeness.

Given a generic linear discrete-time state-space system, the equations of motion can be written as follows:

$$\begin{aligned}\vec{x}_{k+1} &= A\vec{x}_k + B\vec{u}_k \\ \vec{y}_k &= C\vec{x}_k + D\vec{u}_k\end{aligned}\quad (5)$$

Note that in the case of the Atlantis, \vec{u}_k is the rudder slew rate, $\dot{\delta}$, and \vec{y}_k is, as before, $[y \ \psi \ \delta]^T$. That is, the OKID will generate a set of $[A, B, C, D]$ matrices that match the input and output data.

It has been shown that the triplet, $[A, B, C]$ is not unique, but can be transformed through any similarity transform (i.e., 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, is invariant under similarity transforms. That is, for an initial condition of $\vec{x}_0 = 0$, and for the input $\vec{u}_0 = 1$ and $\vec{u}_k = 0$ for all $k > 0$, the system response is:

$$\begin{aligned}Y_0 &= Du_0 \\ Y_1 &= CBu_0 \\ Y_2 &= CABu_0 \\ Y_3 &= CA^2Bu_0 \\ &\vdots \\ Y_k &= CA^{k-1}Bu_0\end{aligned}\quad (6)$$

These Markov parameters, $[Y_0 \ Y_1 \ \dots \ Y_k]$, are clearly unchanged by any similarity transform, and are thus a unique representation of a linear time invariant system.

When these Markov parameters are assembled into a specific form – the generalized Hankel matrix of Equation 7 – this matrix can be decomposed into the Observability matrix, a state transition matrix, and the Controllability matrix (Equation 8); thus the Hankel matrix (in a noise-free case) will always have rank n , where n is the system order.

$$H(k-1) = \begin{bmatrix} Y_k & Y_{k+1} & \dots & Y_{k+\beta-1} \\ Y_{k+1} & Y_{k+2} & \dots & Y_{k+\beta} \\ \vdots & \vdots & \ddots & \vdots \\ Y_{k+\alpha-1} & Y_{k+\alpha} & \dots & Y_{k+\alpha+\beta-2} \end{bmatrix}\quad (7)$$

$$\begin{aligned}H(k-1) &= \begin{bmatrix} C \\ CA \\ CA^2 \\ \vdots \\ CA^{\alpha-1} \end{bmatrix} A^{k-1} [B \ AB \ A^2B \ \dots \ A^{\beta-1}B]\end{aligned}\quad (8)$$

$$H(k-1) = CA^{k-1}\mathcal{O}\quad (9)$$

That is, regardless of how far out one takes the Hankel matrix (via parameters α and β), the rank of the Hankel matrix is always n , due to the fact that Controllability (\mathcal{C}) and Observability (\mathcal{O}) matrices can be at most of Rank n .

$$\Re(\mathcal{C}) \leq n\quad (10)$$

$$\Re(\mathcal{O}) \leq n\quad (11)$$

$$\Re(H(k-1)) \leq n\quad (12)$$

Because noise will corrupt the rank deficiency of the Hankel matrix (that is, for real noisy data, the Hankel matrix will 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 Atlantis. A more complete treatment of the subject can be found in [10].

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 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. To see this, the recursion from Equation 5 is carried out (assuming for simplicity that the initial state, \vec{x}_0 , is 0):

$$\begin{aligned}y_0 &= Du_0 \\ y_1 &= CBu_0 + Du_1 \\ y_2 &= CABu_0 + CBu_1 + Du_2 \\ y_3 &= CA^2Bu_0 + CABu_1 + CBu_2 + Du_3 \\ &\vdots \\ y_k &= CA^{k-1}Bu_0 + CA^{k-2}Bu_1 + \dots + CABu_{k-2} \\ &\quad + CBu_{k-1} + Du_k\end{aligned}\quad (13)$$

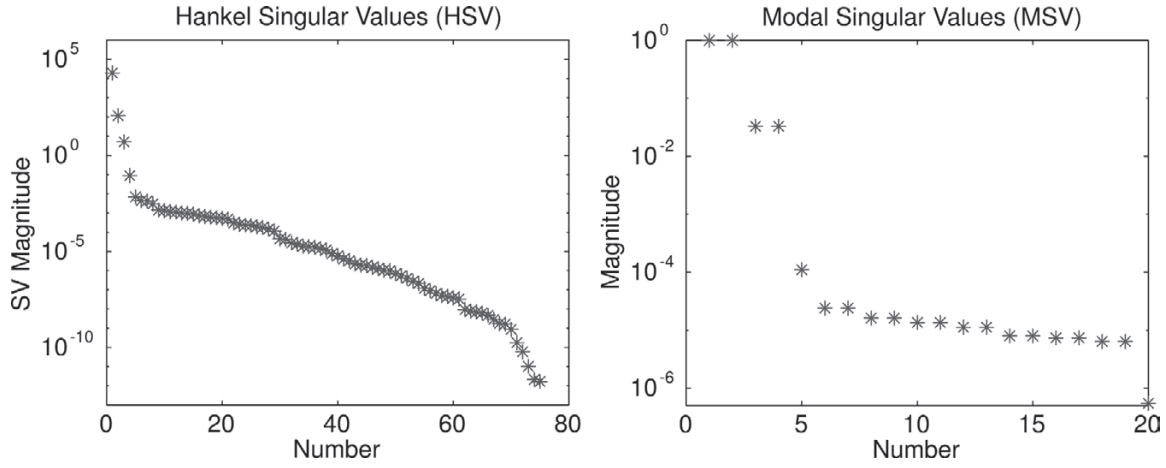


Fig. 11—The plot of the Hankel singular values of the system, showing that the experimental data is best represented by a 4th order system (left), and a singular value plot of the system modes showing again a large drop after 4 modes, again indicating that the system is 4th order (right).

These equations can be gathered into a matrix form (note that the OKID method is general for multiple input multiple output systems, which is to say that both Y_k and u_k can be vectors). These gathered equations result in a topelitz matrix form:

$$\begin{bmatrix} y_k \\ \vdots \\ y_2 \\ y_1 \\ y_0 \end{bmatrix} = \begin{bmatrix} D & CB & CAB & \dots & CA^{k-1}B \\ & \ddots & \ddots & \ddots & \vdots \\ & & D & CB & CAB \\ & & & D & CB \\ & & & & D \end{bmatrix} \begin{bmatrix} u_k \\ \vdots \\ u_2 \\ u_1 \\ u_0 \end{bmatrix} \quad (14)$$

Note that the entries of this matrix are nothing other than the system Markov parameters, in an upper triangular form. Under normal circumstances, there are not enough equations available to solve for all of the Markov parameters. That is, given the input stream (u_k) and the outputs (y_k) for some span of time (k from 0 to r), we cannot simply solve for the $[D, CB, CAB, \dots, CA^k B]$ Markov parameters.

Examination of Equation 14 shows that if for some k , the quantity $A^k = 0$, then the number of unknown Markov parameters can be greatly reduced, or truncated. $A^k = 0$ for some k , however, is the definition of asymptotic stability, which is simply to say that the effect of an input k steps in the past is negligible on the current outputs. 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:

$$\begin{aligned} \bar{x}_{k+1} &= A\bar{x}_k + B\bar{u}_k + G\bar{y}_k - G\bar{y}_k \\ \bar{x}_{k+1} &= [A + GC]\bar{x}_k + [B + GD]\bar{u}_k - G\bar{y}_k \\ \bar{x}_{k+1} &= \bar{A}\bar{x}_k + \bar{B}\bar{u}_k \end{aligned} \quad (15)$$

where:

$$\begin{aligned} \bar{A} &= [A + GC] \\ \bar{B} &= [B + GD - G] \\ \bar{v}_k &= \begin{bmatrix} \bar{u}_k \\ \bar{y}_k \end{bmatrix} \end{aligned} \quad (16)$$

Thus, the system stability can be augmented through an observer, which has the effect of making $\bar{A}^p \simeq 0$ for some p that is sufficiently large. With that assumption, there are enough equations to solve for the Markov parameters established through a least-squares solution [8]. It is useful to note that the realization also provides a pseudo-Kalman observer. The observer orthogonalizes the residuals to 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 [10], was used to identify the sailboat dynamics from the experimental data.

An SVD of the aggregate velocity-normalized data for the Atlantis demonstrated a large drop in the magnitude of the singular values from the fourth to the fifth, indicating a system order, n , of four (Figure 11, left). In addition, modal singular values (Figure 11, right) of all catamaran models of order higher than four exhibited a two order-of-magnitude drop from the fourth modes to modes higher than four. System reconstruction for the identified dynamics also matches well, showing predictive performance that matched within 0.1 m of cross track error, within 1.5 deg of heading error, and within 4 deg for rudder angle. Note that these were open loop tests, without the Kalman

filter added in, which improved prediction greatly. Thus, using these results, the LQG controller was designed, and simulations carried out to validate the controller performance. Once satisfied with these simulations, experimental trials were performed in order to validate the concept.

The implementation of the identified controller and estimator is fairly straightforward. From the OKID methodology, the linear system matrices $[A, B, C, D]$ and the estimator G are extracted. The control gain matrix, K , is extracted from the LQR design (with the cost function in Equation 4). Abstractly, the controller is simply a mathematical function that maps (with memory), a sequence of system outputs to a set of control inputs. In this specific case, the first step is to normalize the crosstrack error and heading error by the velocity (with a lower bound of 1 m/s to prevent noise amplification) as shown in Equation 2.

Given that this implementation includes an estimator, the initial estimate of the state is set to 0 (note that there are several methods for attempting to find a better initial estimate for the state, such as taking the pseudo-inverse of the C matrix and multiplying it by the current measured outputs). The initial control, u , is also set to 0.

At any given time-step, the control is calculated from the current state estimate:

$$u = -[K]\hat{x} \quad (17)$$

and then the estimator is propagated forward one step in time given the current control and measured system outputs:

$$\hat{x}_+ = [A + GC]\hat{x}_- + [B + GD]u - [G]\tilde{y} \quad (18)$$

where \hat{x}_+ is the future state estimate, \hat{x}_- is the previous estimate, u is the actuator command, and \tilde{y} is the speed normalized sensor readings.

Note that the control portion of the propagation $[B + GD]$ could be implicitly brought into the \hat{x}_k term, but in this case was not due to saturation limits imposed by the actuator hardware.

TROLLING MOTOR TESTS

While the wing-sail was still under construction at Cris Hawkins Consulting in Santa Rosa, the system identification and controller tasks had already been completed. At this point, in order to test out the controllers, a MinKota electric trolling motor was used to simulate the presence of the wing-sail and wind. This was done by mounting the trolling motor at the sailboat center of gravity (CG), and turning the trolling motor such that its direction of thrust was canted off the centerline by more than 40 degrees.



Fig. 12—The Atlantis on an unmanned trajectory being controlled by the identified LQG controller. Propulsion is from a MinKota trolling motor running at 12, 24, and 36 Volts. The motor is canted off the center line to simulate off-center thrust as would be seen by the wing-sail.

Since the dynamics of the catamaran are greatly affected both by the velocity through the water and the displacement weight of the hulls, the Atlantis was ballasted with an additional 75 kg of lead ballast (in the form of batteries) to bring the all-up weight of the boat to the same as the weight as it would have had with the wing-sail installed. Also, in order to test the controllers at various speeds, the MinKota trolling motor was run with 12, 24, and 36 V at approximately 65 A. This changes the speed of the boat through the water, simulating changes in wind velocity. In order to simulate changes in wing direction, the MinKota trolling motor was turned through various angles while the controller was regulating the path to a line.

In Figure 12, the Atlantis with the trolling motor can be seen. The trolling motor is at the center of the boat, and the lead batteries provide the ballast. As pictured, the boat was run unmanned, with the GNC computer providing all navigation, at a speed of approximately 2 m/s (4 Kn). Of note is the fact that the anemometer is located at the front wooden crossbeam. This is only a temporary location, and moving the sensor's physical location is very easy due to the CAN bus architecture employed on the Atlantis.

Figure 13 shows a 500 m long typical autonomous pass while under computer control. Note that the computer regulates the path to the line, but that the turn is performed open loop with a feed-forward command. To the scale pictured in Figure 13, the recorded position data shows very little cross track error. This is in spite of the fact that the currents were changing, and the wind and waves were all injecting disturbances into the system. In Figure 14, a close look at the errors in the first part of the path shown in Figure 13

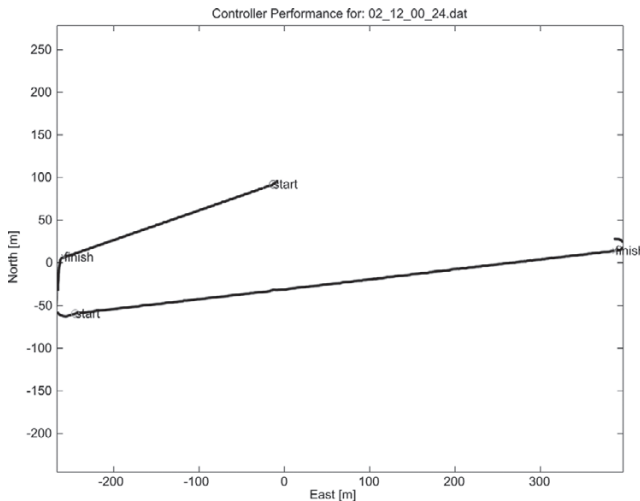


Fig. 13—The trajectory as recorded by the Ag122 Differential GPS receiver while under computer control. Note that feedback control is only used on the straight line segments. Turns are accomplished by open-loop feed-forward commands. To the resolution of this image, cross track errors are less than one pixel wide.




Fig. 14—Close up of the first section pictured in Figure 13, showing very precise control while under trolling motor propulsion and autonomous control. The mean of the path is less than 3 cm., and the standard deviation is less than 10 cm. Note that the presence of a current is indicated by a constant heading error (middle) as well as a mismatch between the GPS velocity and the hull speed sensors (bottom).

reveals that the mean was less than 3 cm, and the standard deviation was less than 10 cm.

Of interest, the azimuth shows a -20 degree bias for most of the path length of the run pictured in Figure 13, which is due to current. This can be verified by looking at the velocity plot at the bottom of Figure 14, where the top line is the hull-speed sensor, and the smooth lower line is GPS velocity. The difference in these two is current, and it can be seen in spite of the high frequency noise of the hull-speed sensor (due to the placement

behind the centerboards). By calculation, the current was 0.62 m/s at an angle of 52 deg to the Atlantis path, coming from the port side of the boat. This is a current speed that is close to 30% of the actual speed of the boat, and can be considered quite a large disturbance.

WING-SAIL TESTS

In order to validate the performance of the controllers and all-up system, closed loop control experiments were performed in Redwood City Harbor, California, on 27 January 2001. These tests were intended to verify that the closed loop controllers were capable of precise line following with the increased disturbances due to the wing-sail propulsion. No modifications were made to the controller design, and the tests were run on a day with approximately 12 kn (or 6 m/s) of wind, with gusts up to the 20 kn (or 10 m/s) range. For a typical pass, the velocity of the Atlantis was again approximately 2 m/s (4 kn), or approximately 50% of the true wind speed. Note that the Atlantis was built for precision, and not for performance. Higher velocities were well within the capability of the wing, but were not attempted.

Experimental data from the anemometer shows the angle of the wind with respect to the wing (angle of attack) to vary ± 20 degrees from nominal. This demonstrates the requirement for a self-trimming wing; without the ability to trim quickly to a new angle of attack, the wing would remain stalled most of the time, and the thrust generated would be minimal. The ability to respond quickly to a new angle of attack by rotating into the new trim condition allows the wing to absorb these transient gusts and continue to provide full thrust with a reduced heeling moment.

Qualitatively, the wing-sail performed even better than anticipated. With the tail centered, there was no tendency for the Atlantis to heel whatsoever, and the absence of aeroelastic instability (sail luffing) made the entire event very quiet. Upon turning the trailing edge of the tail in the direction of desired travel, the Atlantis smoothly accelerated to speed and quietly continued on her course. Even large gusts simply caused the Atlantis' wing to quickly stall and, with only a slight shudder, reposition to a new angle of attack (as evidenced by the yarn tufts on the wing surface).

More impressive was the ability to sail pointed very high into the wind. Upon analyzing the data, it was demonstrated that the Atlantis was capable of sailing to within 25 deg of the true wind direction. At one point, a conventional sailboat came about behind the Atlantis and started luffing a full 15 degrees off the wind from where the Atlantis was making headway. This is clearly a result of




Fig. 15—Satellite photograph of harbor where the Atlantis was sailed under computer control. The white dots are data recorded from a human sailor using a conventional sloop rig. The black dots show the data from various closed loop computer controlled segments. Note how straight and “unnatural” they are. Simply, they do not look human.

the improved aerodynamics of the rigid wing, and a vindication of the self-trimming arrangement over a conventional sail. While further experimental studies are required to quantitatively measure the performance increase of the wing-sail, current results indicate very promising discoveries ahead.

Figure 15 shows a satellite picture of the harbor where both the trolling motor and wing-sail tests were performed. The white dots are from a previous year, when the Atlantis was conventionally sailed with a sloop rig, and was sailed by a human captain. The black dots indicate the various closed loop control passes from the recent tests. Note that the white trace has a curving, “human,” look to it, whereas the black trace looks like a ruler was placed upon the photograph and a line drawn. Qualitatively, the computer control simply does not look like a human pilot.

Figure 16 is, once again, a closer look at a bird’s eye view of a set of computer controlled traces. The control system regulated about the lines in between each “start” and “end” pair, and the turns in between were performed open-loop in a feed-forward sense. Figure 17 presents a close-up of the first path of the regulated control, and looks at the cross track error, azimuth error, and velocities. Note that the dark line in the top of the velocity graph (the bottom panel of Figure 17) is the wind speed, and can be seen to vary well over 50% of nominal.

The mean of the cross track error is less than 3 cm, and the standard deviation is less than 30 cm. Note that this is the Sailboat Technical

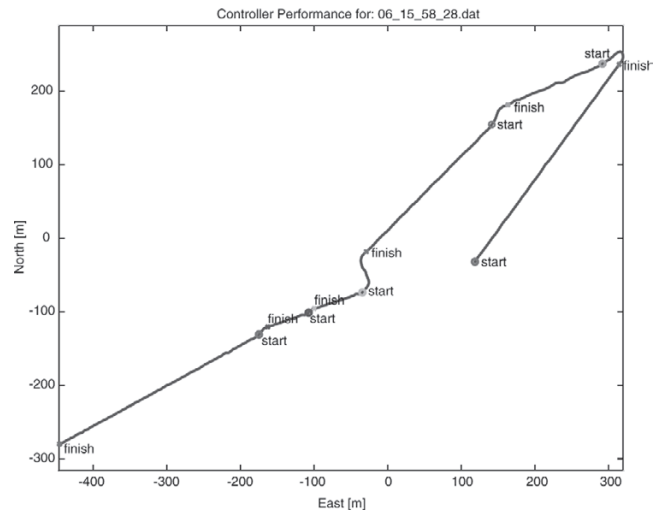


Fig. 16—Birds eye view of Atlantis under computer control, propelled by the wing-sail. Path regulation happens between each “start” and “end” pair, with the curves in between being performed open-loop with a feed-forward command.

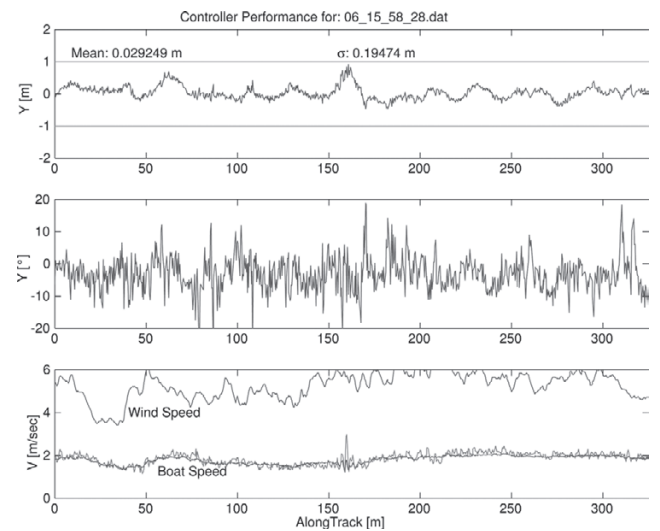


Fig. 17—Close up of one of the control segments displayed in Figure 16, showing a mean of less than 3 cm., and a standard deviation of less than 20 cm. Note that in this case, there was no systematic current, though the wind can be seen to vary more than 50% of nominal on the top line of the bottom plot.

Error (STE), the sailing analog of Flight Technical Error, that is, the difference between the measured position and the reference position. Previous characterization of the coast-guard differential GPS receiver indicated that the Navigation Sensor Error (NSE) is approximately 36 cm, thus the Total System Error (TSE) is less than 1 m [3].

Figure 18 presents the aggregate of all controlled sailing runs overlaid one upon the other. Along with bounds indicating ± 1 m, the differences in path length have to do with the location of the shore, and the desire not to run aground. Depending on the path chosen, longer or shorter distances were traversed. At no time does the controlled

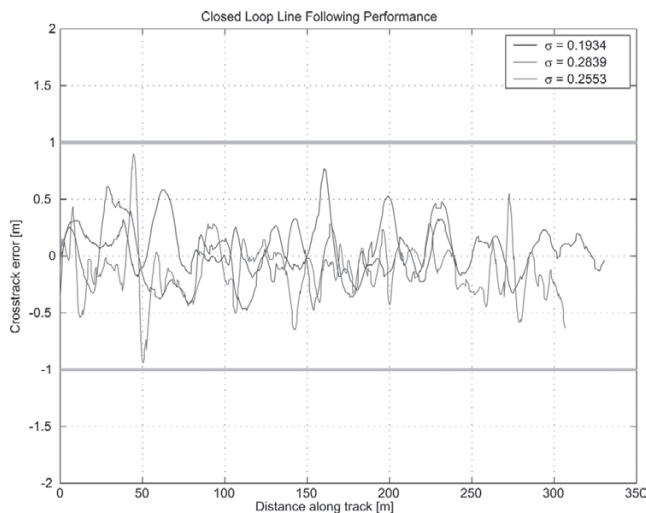


Fig. 18—Aggregate plot of computer controlled sailing passes, with lines at ± 1 meter bounds, overlaid on top of one another. The differences in path length have to do with distance to shore, and the desire not to run aground. The controller keeps the Atlantis well within the one meter bound, and shows a standard deviation of less than 30 cm.

performance of the system exceed the one meter bound. As a basis for comparison, the specifications for the top-of-the-line AutoHelm autopilot indicate a cross track accuracy of 0.05 nmi, or 92.6 m.

CONCLUSION

It has been demonstrated that with the combined advances in GPS technology, and the advent of low cost sensors, an unmanned sailboat can be built that can navigate with unprecedented levels of accuracy. By utilizing a novel wing-sail propulsion system, the difficulties of actuating a sail have been overcome, and high authority control can be realized. A demonstrated Sailboat Technical Error (STE) in line following of less than 0.3 m ($1 - \sigma$) was achieved in challenging conditions. Combined with a Navigation Sensor Error (NSE) of 0.36 m, this yields a Total System Error (TSE) of less than 1 m.

FUTURE WORK

While this project represents a great deal of progress on the concept of unmanned sailing vessels, there remains much more to do in order to make the Atlantis more than an academic research platform. Control has been demonstrated on simple straight line segments. However, more complicated trajectories are required than simple lines. Furthermore, these segments must be linked together in some smooth manner in order to create a viable mission for an autonomous sailing vessel.

While tacking and jibing are very simple with a rigid self-trimming wing-sail, trajectory generation

must occur whenever the desired waypoint is unreachable due to wind direction, which includes tacking and jibing when necessary.

Lastly, due to the nature of automatic control, rearward sailing is only slightly more difficult than forward sailing. With this ability, station keeping becomes a viable maneuver, and one that should be most useful.

We are currently working with the Atlantis at UC Santa Cruz to make the Atlantis more robust, with improved sensors, better integration, and to extend her capabilities to further demonstrate the viability of this craft.

ACKNOWLEDGMENTS

The work in this paper was the result of research sponsored by University of California, Santa Cruz. The initial part of this research was conducted under a birdseed grant from the Stanford Office of Technology Licensing (OTL). The author gratefully acknowledges both the OTL and the Principal Investigator (and Thesis Advisor), Dr. Bradford Parkinson, of Stanford University for support and sharing the data.

REFERENCES

1. Wolfhard, L., *CAN Sytem Engineering: From Theory to Practical Application*.
2. Wahba, Grace, Problem 65-1 (Solution), *SIAM, Review*, 8: 384–386, 1966.
3. Elkaim, G. H., *System Identification for Precision Control of a WingSailed GPS-Guided Catamaran*, PhD thesis, Stanford University, Stanford, CA, 2001.
4. Gebre-Egziabher, D., Elkaim, G. H., Powell, J. D., and Parkinson, B. W., *A Gyro-Free, Quaternion Based Attitude Determination System Suitable for Implementation Using Low-Cost Sensors*, Proceedings of the IEEE Position Location and Navigation Symposium, PLANS 2000, IEE, 2000, pp. 185–192.
5. Gebre-Egziabher, D., and Elkaim, G. H., *MAV Attitude Determination from Observations of Earth's Magnetic and Gravity Field Vectors*, AIAA Aerospace Electronic Systems Journal, submitted for publication.
6. Fossen, T. I., *Guidance and Control of Ocean Vehicles*.
7. Elkaim, G. H., O'Connor, M. L., and Parkinson, W. B., *System Identification and Robust Control of a GPS-Guided Farm Tractor*, Proceedings of the Institute of Navigation ION-GPS Conference, ION, 1997, pp. 1415–1426.
8. Juang, J.-N., *Applied System Identification*.
9. Juang, J.-N., Cooper, J. E., and Wright, J. R., *An Eigensystem Realization Algorithm Using Data Correlations (ERA/DC) for Modal Parameter Identification*, Control-Theory and Advanced Technology, 4(1), 1988.
10. 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, 117(2), 1995.