JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 27, No. 3, May–June 2004

Nanosatellite Attitude Stabilization Using Passive Aerodynamics and Active Magnetic Torquing

Mark L. Psiaki*
Cornell University, Ithaca, New York 14853-7501

Aerodynamic and magnetic torques have been used to develop a new three-axis attitude stabilization system for a nadir-pointing low-Earth-orbiting spacecraft. The goal is to develop a crude three-axis stabilization system for a 1-kg nanosatellite and to stay within a mass budget of 0.1 kg and an average power budget of 0.06 W. The system uses passive aerodynamic drag torques to stabilize pitch and yaw. The aerodynamic design resembles a badminton shuttlecock and uses thin metallic "feathers" that deploy from the rear of a small cubic main spacecraft bus. The active magnetic control system stabilizes roll and damps pitch and yaw by feeding magnetometer measurements back to magnetic torque coils. It also includes a feedforward term that is based on the expected magnetic field vector direction in spacecraft coordinates when the spacecraft is in its commanded orientation. The system has been tested via simulation. Global stabilization is achieved for all altitudes below 500 km, as evidenced by the system's ability to capture the nadir-pointing attitude when starting from a tumble. Worst-case pointing errors are below 25 deg per axis at altitudes below 450 km, and these errors can be decreased by a decrease in orbital inclination, asymmetries, or both.

I. Introduction

MANY Earth-orbiting spacecraft have a nadir-pointing requirement that requires full three-axis attitude stabilization. Typical nadir-pointing stabilization systems use passive gravity gradient booms, semi-active momentum bias systems that include a pitch momentum wheel and magnetic torquers, or fully active systems that include a suite of reaction wheels, control moment gyros, or thrusters to implement full three-axis control.

The present paper deals with the design of a three-axis stabilization system for a nadir-pointing 1-kg nanosatellite. This satellite must fit within a 0.1-m cube during launch. It has a continuous power budget of 1.5 W, which is supplied by a power system that uses batteries and solar cells mounted on the outer surface of the cube. It must operate at altitudes in the range 300–500 km. It requires pitch and roll stabilization in order to orient a global positioning system (GPS) patch antenna toward zenith and a communications antenna toward nadir. The pointing requirements are loose; the system can tolerate roll and pitch errors on the order of 20 deg. The final stabilization system design achieves a similar level of yaw pointing accuracy despite the lack of a yaw requirement.

Traditional attitude stabilization systems are impractical for this satellite. Reaction wheels and momentum wheels do not fit within its limited weight and power budget. Magnetic torque coils can be used, but it is difficult to guarantee global stability with only magnetic torque because of the inability to torque about the local magnetic field direction. The design does not require orbital maneuvering capability and, therefore, does not include thrusters. At altitudes below 400 km, aerodynamic drag torque tends to overwhelm the gravity gradient torque for practical lightweight deployable boom designs.

This paper's new aerodynamic design turns the liability of a large drag torque into an asset. Preliminary design studies of a gravity-gradient system indicated that the aerodynamic torques turn the gravity-gradient boom into the tail of an arrow. This arrow concept has been modified to become a badminton shuttlecock-type design for stabilization of the yaw and pitch axes. A shuttlecock is a better analogy for the aerodynamics of the new system because an arrow relies on lift forces whereas a shuttlecock relies on drag. Lift is very small compared to drag at the high Knudsen numbers (ratio of the mean free path of an air molecule to a reference dimension of the spacecraft) that typify low-altitude space flight.

The new system is similar to a passive aerodynamic stabilization system that is described in Refs. 1–3. That system's goal is to stabilize only yaw and pitch by the use of a totally passive system operating at altitudes of 250–325 km and below. The design resembles a long stove pipe with the center of mass biased toward one end. It tends to fly with its axis of rotational symmetry aligned with the air-relative velocity vector and with the center of mass oriented forward. The design includes magnetic hysteresis damping rods for passive damping. Its roll axis is not stabilized. A flight test was conducted off of space shuttle mission STS-77 in May 1996. Video recordings of the system's performance indicate that it achieved pitch and yaw stabilization after several days of operation, despite the fact that it tumbled at a rate of more than 2 deg/s on release from the space shuttle.

The new design stabilizes all three axes by combining the pitch–yaw stabilization concept of Refs. 1–3 with an active magnetic feedback control system. The new aerodynamic pitch–yaw stabilization design is more practical for a small satellite. It replaces the long solid cylinder of Refs. 1–3 with a compact cubic main bus from which, to use the shuttlecock analogy, three small "feathers" deploy after launch. These feathers move the aerodynamic center of pressure aft of the center of mass to produce the desired stability.

The active magnetic stabilization control law is similar to that of Ref. 4. It feeds three-axis magnetometer measurements back to the magnetic torque coils. The measured magnetic field direction vector is compared with the value that it would take on if the satellite were pointed toward nadir. Errors are fed back with a special proportional–integral–derivative control law. A similar system has been shown to stabilize pitch and to add roll–yaw damping when a momentum wheel is present for roll–yaw stiffening.[4] In the present case, the magnetic feedback system adds damping to the pitch and yaw axes, and it stabilizes roll. This method of damping pitch and yaw enables the design to overcome problems of classical and parametric resonance that can occur when magnetic hysteresis rods provide the only damping.[2]

The present paper makes several contributions to the technology of aerodynamically stabilized spacecraft. First, it presents an aerodynamic model that includes more physics than the one given in Ref. 1. Second, it presents and analyzes a different aerodynamic design for

Received 18 April 2003; revision received 30 September 2003; accepted for publication 19 October 2003.

*Associate Professor, Sibley School of Mechanical and Aerospace Engineering. Associate Fellow AIAA.

the pitch–yaw stabilization system, one based on a shuttlecock-type configuration rather than a stove pipe configuration. Third, it shows how to damp pitch and yaw actively and stabilize roll by the use of a relatively simple, robust magnetometer-to-magnetic-torquer feedback controller. Fourth, it considers the effects on the pointing accuracy of off-nominal parameters. These include lift-induced aerodynamic roll torques due to twist of the shuttlecock's feathers and aerodynamic pitch and yaw disturbances due to static mass imbalance.

The design, analysis, and simulation testing of this new system are discussed in Secs. II–V of this paper. Section II presents the aerodynamic pitch–yaw stabilization design and two models of it. Section III defines the magnetic feedback controller. Section IV describes the simulation studies that have been conducted, and Sec. V presents results from these studies. Sec. VI presents the paper's conclusions.

II. Design and Modeling of the Aerodynamic Pitch–Yaw Stabilization System

A. Aerodynamic Design

The pitch–yaw stabilization system uses aerodynamic drag to develop pitch and yaw stiffness. The basic idea is to place the center of drag pressure behind the center of mass so that the system works somewhat like a badminton shuttlecock (Fig. 1). Each deployable feather shown in Fig. 1 can be made with a thin bowed strip of steel or Kevlar that is coiled when in its stowed configuration, much like a retractable carpenter's tape measure. A prototype feather system has been designed, built, and successfully deployed.

This design uses differential drag on its three feathers to stabilize pitch and yaw. Suppose that the system is oriented with a positive pitch angle (nose up). Then the feathers, which are rear-facing, tilt downward. This decreases the two top feathers' projected areas perpendicular to the flow, thus reducing their drag, and it decreases their vertical moment arms to the center of mass. Conversely, the lower feather has an increased projected area perpendicular to the flow, which causes increased drag, and an increased vertical moment arm. If the center of mass of the system is near the center of the main bus, then these effects combine to produce a restoring nosedown drag torque. A similar situation holds for rotations about the yaw axis.

B. Simplified Aerodynamic Pitch–Yaw Stiffness Model

A simplified aerodynamic pitch–yaw stiffness model demonstrates the fundamentals of why this system works. Consider the four-feathered design shown in Fig. 2. The geometric quantities defined in Fig. 2 are α, the aerodynamic angle of attack; γ, the angle between each feather and the x axis; l, the length of each feather;

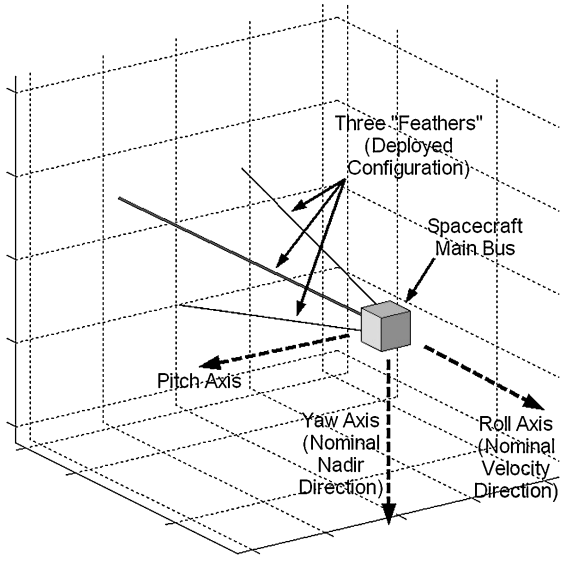

Fig. 1 Three-dimensional view of the shuttlecock aerodynamic pitch–yaw stabilization system.

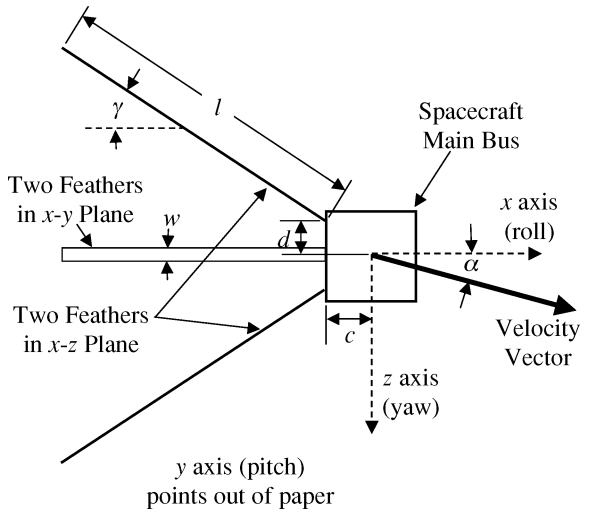

Fig. 2 Geometry of a simple model of aerodynamic pitch stabilization of a four-feathered design.

d, the distance from the in-board tip of each feather to the x axis; c, the distance of the in-board tip of each feather behind the z axis; and w, the width of each feather. Simple geometry can be used to show that the areas of the two x–z-plane feathers projected perpendicular to the velocity vector are

$$A_{\text{up}} = lw\sin(\gamma - \alpha) \tag{1a}$$

$$A_{\text{lo}} = lw\sin(\gamma + \alpha) \tag{1b}$$

where A_{up} corresponds to the upper feather and A_{lo} corresponds to the lower feather. The drag force on each feather acts through its geometric center in a direction opposite to the velocity vector. The projections of the moment arms of the two x–z-plane feathers perpendicular to the velocity vector are

$$h_{\text{up}} = -c\sin\alpha + d\cos\alpha + (l/2)\sin(\gamma - \alpha) \tag{2a}$$

$$h_{\text{lo}} = -c\sin\alpha - d\cos\alpha - (l/2)\sin(\gamma + \alpha) \tag{2b}$$

These projected moment arms are measured from the origin of the coordinate system at the center of the spacecraft bus, which is presumed to be the system's center of mass. They are positive if the feather's geometric center is above the line of the velocity vector.

The two x–y-plane feathers are angled into and out of the plane of the paper in Fig. 2. Each x–y feather has a projected area perpendicular to the velocity vector equal to

$$A_{xy} = lw\sin\gamma\cos\alpha \tag{3}$$

and a projected moment arm to the center of mass equal to

$$h_{xy} = -[c + (l/2)\cos\gamma]\sin\alpha \tag{4}$$

The geometric formulas in Eqs. (1a–4) can be used in conjunction with a crude aerodynamic model to calculate the aerodynamic pitch torque. This approximate model ignores lift and uses a drag coefficient of two for all surfaces that are impacted by the flow. The reference area of each surface equals its projected area perpendicular to the velocity vector. If surfaces other than the feathers are neglected, then this analysis yields the following pitch torque model:

$$\begin{aligned} n_{y\,\text{aero}} &= \bar{q}2(h_{\text{up}}A_{\text{up}} + h_{\text{lo}}A_{\text{lo}} + 2h_{xy}A_{xy}) \\ &= -\bar{q}2lw\sin(2\alpha)[2c\sin\gamma + d\cos\gamma + (3l/4)\sin(2\gamma)] \end{aligned} \tag{5}$$

where $\bar{q} = 0.5\,\rho V^2$ is the dynamic pressure. The linearized pitch stiffness is

$$K = -\left.\frac{\mathrm{d}n_{y\,\text{aero}}}{\mathrm{d}\alpha}\right|_{\alpha=0} = \bar{q}4lw\left[2c\sin\gamma + d\cos\gamma + \frac{3l}{4}\sin(2\gamma)\right] \tag{6}$$

The aerodynamic stiffness in Eq. (6) is surprisingly accurate. It has been checked against the next section's model, which has more physics in it. For a representative case, Eq. (6) yields a stiffness that equals 88% of the value from the more precise model.

Similar crude models could be developed for other numbers of feathers. The four-feathered model has been chosen for this example because it simplifies the calculations.

C. Torque Model Based on Free-Molecular Aerodynamics

Aerodynamic design and simulation calculations have been carried out with the free-molecular aerodynamic force model of Ref. 5. This model gives the normal pressure and the shearing stress on a surface of the spacecraft as functions of the dynamic pressure, the angle of attack, and various thermodynamic quantities. Consider the geometry of Fig. 3. The quantities in Fig. 3 are $\boldsymbol{v}$, the velocity vector of the surface element with respect to the atmosphere; $\hat{\boldsymbol{u}}$, the outward-pointing unit normal vector; α, the angle of attack; p, the total pressure; and τ, the shearing stress. The heavy vectors next to the p and τ symbols show the directions in which they apply force to the surface.

Formulas for the total pressure and the shearing stress are given in Ref. 5. Equivalent formulas are

$$\frac{p}{\bar{q}} = \left\{\left[\frac{2-\sigma_n}{\sqrt{\pi}}\right]\sin\alpha + \frac{\sigma_n}{2s}\sqrt{\frac{T_s}{T_a}}\right\} \times \left\{\frac{1}{s}e^{-s^2\sin^2\alpha} + \sqrt{\pi}\,[1+\operatorname{erf}(s\sin\alpha)]\sin\alpha\right\} + \left[\frac{2-\sigma_n}{2s^2}\right][1+\operatorname{erf}(s\sin\alpha)] \tag{7a}$$

$$\frac{\tau}{\bar{q}\cos\alpha} = \sigma_t\left\{\frac{1}{s\sqrt{\pi}}e^{-s^2\sin^2\alpha} + [1+\operatorname{erf}(s\sin\alpha)]\sin\alpha\right\} \tag{7b}$$

where p is the pressure, τ is the shearing stress, σ_n and σ_t are the normal and tangential accommodation coefficients, T_s is the absolute temperature of the spacecraft surface, T_a is the atmospheric temperature, erf( ) is the error function

$$\operatorname{erf}(\beta) = \frac{2}{\sqrt{\pi}}\int_0^{\beta} e^{-\eta^2}\,\mathrm{d}\eta \tag{8}$$

and s is the airspeed, nondimensionalized by the mean molecular speed of the atmosphere

$$s = \sqrt{\frac{M\|\boldsymbol{v}\|^2}{\mathbf{2}RT_a}} \tag{9}$$

The quantities M and R in the formula for s are, respectively, the mean molar mass of the atmosphere and the universal gas constant. Note also that $\sin\alpha = \hat{\boldsymbol{u}}^T\hat{\boldsymbol{v}}$, where $\hat{\boldsymbol{v}} = \boldsymbol{v}/\|\boldsymbol{v}\|$.

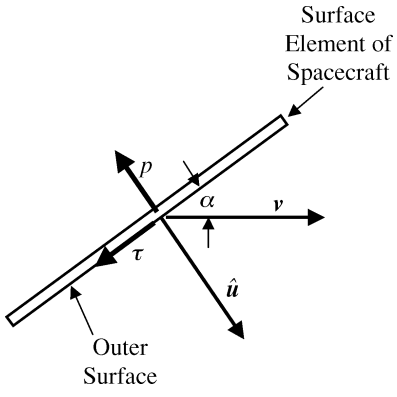

Fig. 3 Geometry of free-molecular aerodynamic force model.

The accommodation coefficients σ_n and σ_t model the fact that some impinging air molecules come to thermal equilibrium with the spacecraft surface and then get reemitted thermally, whereas others reflect specularly. Specular reflection means that the angle of incidence equals the angle of reflection and that the relative speed remains unchanged. The values $\sigma_n = \sigma_t = 1$ correspond to full thermal accommodation of all of the impinging molecules, and $\sigma_n = \sigma_t = 0$ corresponds to specular reflection of all of the air molecules. Typical values of both coefficients fall in the range 0.85–1.0, which means that most of the air molecules get reemitted thermally. There are no reliable data for these coefficients. It is recommended that one evaluate a given design's performance over a range of possible σ_n and σ_t values or that one use worst-case values.

The pressure and shearing stress can be used in conjunction with the vectors $\hat{\boldsymbol{u}}$ and $\hat{\boldsymbol{v}}$ to compute force and torque. The force on a flat surface of area δA is

$$\delta\boldsymbol{f} = \delta A[-\hat{\boldsymbol{u}}p + (\hat{\boldsymbol{u}}\sin\alpha - \hat{\boldsymbol{v}})(\tau/\cos\alpha)] \tag{10}$$

Note that $(\tau/\cos\alpha)$ is finite even when $\alpha = 90$ deg, per Eq. (7b), and that the vector $(\hat{\boldsymbol{u}}\sin\alpha - \hat{\boldsymbol{v}})$ is the projection of $-\hat{\boldsymbol{v}}$ perpendicular to $\hat{\boldsymbol{u}}$. The net torque on the spacecraft can be computed by summing moments over all external surfaces. It is

$$\boldsymbol{n}_{\text{aero}} = \sum_{i=1}^{N_{\text{sc}}} \boldsymbol{r}_i \times \delta\boldsymbol{f}_i \tag{11}$$

where i is an index of the N_{sc} outer surfaces of the spacecraft, $\boldsymbol{r}_i$ is the position vector of the geometric center of the ith surface measured relative to the spacecraft's center of mass, and $\delta\boldsymbol{f}_i$ is the force on the ith surface from Eq. (10). This study employs only rectangular surfaces.

The atmospheric temperature T_a and the mean molar mass of the atmosphere M are derived as functions of altitude from an atmospheric model. The 1976 U.S. Standard Atmosphere is used.[6] This model does not capture diurnal perturbations or other similar effects. Fortunately, this neglect of perturbations does not have a significant impact at the high values of the non-dimensional velocity s that are typical of space flight.

This aerodynamic model can be used for the spacecraft design considered in Refs. 1 and 2 without the need to consider shading. In this aerodynamic context, shading refers to the situation where one aerodynamic surface lies directly upstream of another aerodynamic surface. Shading does not need to be considered for the design of Refs. 1 and 2 because its outer surface is convex and because the formulas for p and $(\tau/\cos\alpha)$ in Eqs. (7a) and (7b) are valid for negative angles of attack. At high s values, these two functions fall off rapidly to zero as the angle of attack becomes negative.

The aerodynamic model of the shuttlecock design of Fig. 1 uses Eqs. (7a) and (7b), despite the fact that its geometry is not convex. Normally, the model of a nonconvex design must consider the effects of shading and reincidence of emitted air molecules. Both effects can perturb the incident velocity distributions that have been used to derive Eqs. (7a) and (7b).

Shading can be neglected in the present case because the only significant nonconvex shading occurs when a feather lies in the wake of the main spacecraft bus. The feather's effect in the unshaded model is small in this situation. (See Eqs. (1a) and (1b) and Fig. 2.) Shading only serves to reduce this small effect. Thus, the neglect of shading has little impact.

Reincidence of emitted molecules can be neglected for a different reason: Each surface that lies in the half-space that is outside another surface subtends a relatively small solid angle at most points on that other surface because of the nature of the shuttlecock design. The long thin feathers do not subtend much of the field-of-view of any other surface because of their thinness. Each side of the cubic main bus appears small from most of the points on the feathers because, relative to its size, each side is remote from most of each feather.

The aerodynamic model calculations have been implemented in an efficient manner that resembles the technique used in Ref. 1. The model uses tabulated values of the torque vector divided by the

dynamic pressure. These values are given in units of length cubed. They are tabulated on a two-dimensional azimuth-elevation grid of possible spacecraft velocity vector orientations relative to spacecraft coordinates. The nominal grid spacing for most cases considered in this paper is 10 deg. The dynamic simulation of the spacecraft uses the simulated attitude and the simulated orbital velocity and altitude to compute the dynamic pressure and the velocity direction in spacecraft coordinates. The velocity direction is then used to determine the partially dimensional torque vector via bilinear interpolation in the table, and the result is multiplied by the dynamic pressure to complete the aerodynamic torque calculation.

The partially dimensional torque tabulation uses an assumed spacecraft altitude and an assumed mean airspeed in the evaluations of Eqs. (7a), (7b), and (9). This precomputation approach is reasonable because the small deviations from the assumed altitude and airspeed that are typical for space flight in a nearly circular orbit do not have a significant impact on the tabulated ratio of the torque to the dynamic pressure.

The aerodynamic torque model does not include attitude-rate effects. Such effects depend on the nondimensional attitude rate $r\omega/\|\boldsymbol{v}\|$, where r is a characteristic length of the spacecraft. This nondimensional rate is typically very small because of the large magnitude of $\boldsymbol{v}$. Therefore, rotation rate effects need not be included in the aerodynamic model.

D. Example Aerodynamic Torque Models

Torque models have been calculated for several specific configurations. One is the baseline stove pipe configuration used in Ref. 2 for the Gravity And Magnetic Earth Surveyor (GAMES) mission. The free-molecular flow model has been used to calculate its pitch–yaw stiffness at its initial altitude of 325 km. The calculated aerodynamic stiffness, which equals the derivative of the pitch or yaw torque component with respect to the pitch or yaw attitude angle, is 1.6 times larger than that of Ref. 2. This is reasonably close given the differences between the Ref. 2 model and the model of Sec. II.C.

The baseline shuttlecock design considered in the present study has the following parameters. The main bus is a cube, as in Fig. 1, with the length of each side equal to 0.1 m. It has three feathers. Each feather is $l = 1$ m long and $w = 0.00635$ m wide and is canted with respect to the roll axis by $\gamma = 12$ deg. Each feather's base is mounted at $c = 0.05$ m behind the center of the cube and is offset from the spacecraft's x axis by $d = 0.035$ m.

Aerodynamic torque tables have been calculated for various representative orbits. Figure 4 shows the yaw torque divided by the dynamic pressure as a function of the azimuth and elevation of the aerodynamic velocity vector. Figure 4 assumes a circular orbit with an altitude of 400 km and an inclination of 0 deg. It also assumes that the spacecraft surface temperature is $T_s = 300$ K and that the thermal accommodation coefficients are $\sigma_n = 0.85$ and $\sigma_t = 0.90$. Figure 4 shows the yaw stability of the configuration, as is evidenced by the steep climb of the yaw torque as the aerodynamic azimuth angle increases from zero. The yaw stiffness is proportional to the slope of this part of the surface. This stiffness yields a yaw-mode oscillation frequency of 0.00754 rad/s for a typical yaw moment of inertia. This frequency is 6.65 times the orbital frequency.

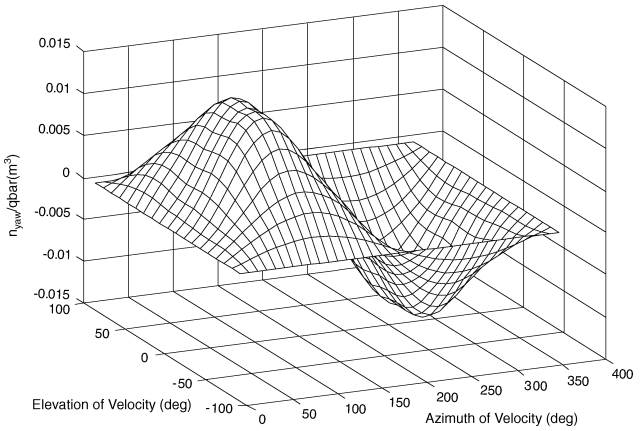

Fig. 4 Tabulated yaw torque divided by dynamic pressure as a function of azimuth and elevation of air-relative velocity measured with respect to spacecraft coordinates.

Another important consideration is the drag coefficient of the design. The design associated with Fig. 4 has a nominal drag coefficient of 3.47 if the aerodynamic reference area equals the area of a face of the main bus, $A_{\text{ref}} = 0.01$ m^2. This is larger than the typical drag coefficient for a spacecraft, $C_D = 2$. The feathers increase the drag coefficient by adding surface area that has a significant nominal angle of attack. The drag coefficient 3.47 gives this 1-kg spacecraft a ballistic coefficient of $m/(C_D A_{\text{ref}}) = 28.8$ kg/m^2. Starting from a circular orbit at 400-km altitude, this translates into an orbital lifetime of only 52 days. An increase of the initial altitude to 500 km would lengthen the satellite's lifetime to 258 days.

This design is slightly more efficient than one like the GAMES design of Ref. 2. A GAMES-like design could achieve an equivalent pitch and yaw stiffness by shaping the spacecraft to be an elongated rectangular parallelepiped. This design would use the same spacecraft bus and would deploy a lightweight extension out of its back that had the same 0.1-m^2 cross-section. The total length of this design would have to be 0.42 m to achieve the same pitch and yaw stiffness if the center of mass were at the center of the leading 0.1-m cube. The corresponding nominal drag coefficient would be 3.56. The 2.6% drag increase of this design vs the birdie design is the result of increased skin-friction drag.

The aerodynamic model of Sec. II.C has also been used to study the effects of off-nominal conditions. Two significant effects that can be studied with this model are pitch–yaw imbalance and roll torques due to twist of the feathers.

If the spacecraft's center of mass is not located on the x axis, then there will be nonzero pitch or yaw torques when the x axis is aligned with the velocity vector. One goal of the aerodynamic design is to minimize the resulting nonzero equilibrium pitch and yaw angles. The nominal shuttlecock design parameters given in this section do a reasonable job of counteracting the effects of imbalance. If the center of mass of the main bus is 0.035 m away from the x axis, which represents an extremely large imbalance, then the equilibrium pitch–yaw bias is only 10.4 deg.

Collective aerodynamic twist of the shuttlecock's feathers causes a roll torque. Twist is defined as rotation about the long axis of a feather. The feathers are thin plates, and collective twist will give rise to small lift forces that will cause a roll torque when in the nominal orientation. This roll torque will result in a constantly increasing roll rate in operation without active stabilization. The aerodynamic roll torque model has been used to examine whether the active magnetic roll stabilization system is capable of counteracting the roll torque that would arise if the collective twist of the feathers equaled the worst-case value. When this effect is studied, it is also necessary to consider the worst-case thermal accommodation coefficients. The values $\sigma_n = 0.85$ and $\sigma_t = 0.90$, that is, the values used to produce Figure 4, approximate the worst case by allowing increased specular reflection and, therefore, increased lift.

III. Active Stabilization Using Magnetic Sensing and Actuation

A nadir-pointing stabilization system can be completed by addition of a pitch–yaw damping system and a roll stabilization system. The goal is to add these functions with a simple, light-weight, low-power active feedback system. The chosen system feeds back magnetometer measurements to magnetic torque coils.

A. Review of the Compass Control Law

The magnetic stabilization system is an enhanced version of the compass control system mentioned in Ref. 4. The most primitive version of the compass control law takes the form

$$\boldsymbol{m} = K_p[\hat{\boldsymbol{b}}_{\text{des}} + \tau_1(\dot{\hat{\boldsymbol{b}}}_{\text{des}} - \dot{\hat{\boldsymbol{b}}})] \tag{12}$$

where $\boldsymbol{m}$ is the commanded magnetic dipole moment vector of the torque coils, $\hat{\boldsymbol{b}}_{\text{des}}$ is the desired magnetic field unit direction vector, $\hat{\boldsymbol{b}}$ is the actual measured magnetic field unit direction vector, K_p is the scalar proportional gain, and τ_1 is the scalar prediction time. The latter two quantities are tuning parameters of the control law. The three vectors in Eq. (12) are given in spacecraft coordinates. The vector $\hat{\boldsymbol{b}}_{\text{des}}$is the direction in which the magnetic field would point in spacecraft coordinates if the spacecraft were oriented in the desired attitude. This vector gets computed by the use of a spherical harmonic expansion of the Earth's magnetic field, the known spacecraft position along its orbit, and the known desired spacecraft attitude.

The control law in Eq. (12) acts like a compass with a damper. The leading term in the bracketed expression points the magnetic dipole in the direction that the magnetic field would point if the attitude were correct. This term pulls the spacecraft attitude toward the desired attitude in the same way that a compass needle gets pulled to align with a magnetic field. The second term in the bracketed expression of Eq. (12) is the derivative term. It provides damping. It is similar to a traditional "bdot" control law, except that it includes the reference time derivative term $\dot{\hat{\boldsymbol{b}}}_{\text{des}}$.

If the spacecraft includes a passive stabilization system that acts around $\hat{\boldsymbol{b}}$, then it tends to stabilize rotation around $\hat{\boldsymbol{b}}$, and the active magnetic controller tends to stabilize the other two axes. This stabilization of the other two axes occurs because the Eq. (12) controller drives $\hat{\boldsymbol{b}}$ to align itself with $\hat{\boldsymbol{b}}_{\text{des}}$. Once this alignment has been achieved, the only possible pointing error is a rotation about $\hat{\boldsymbol{b}}$.

There are various ways to stabilize rotation passively around $\hat{\boldsymbol{b}}$. The design in Ref. 4 uses a pitch momentum wheel to stabilize roll and yaw passively. The current design uses the shuttle-cock's aerodynamics to stabilize pitch and yaw passively. Both designs stabilize rotation around the magnetic field, despite the fact that the magnetic field direction rotates with respect to the desired spacecraft coordinate system as the spacecraft moves along its orbit. In both cases, the spacecraft experiences zero torque from the passive stabilization system and zero torque from the active stabilization system if it is in the correct attitude, but it experiences a restoring torque from the passive stabilization system and restoring and damping torques from the active stabilization system if it is in the wrong attitude.

B. Compass-like Proportional–Integral–Derivative Control Law

The control law of Eq. (12) has been modified in several important ways to develop a magnetic roll stabilization and pitch–yaw damping system for an aerodynamically stabilized shuttlecock spacecraft. The first modification improves the control law's efficiency. The control law in Eq. (12) leaves the magnetic torque coils turned on at $\boldsymbol{m} = K_p\hat{\boldsymbol{b}}$ after the spacecraft reaches the desired attitude, even though this $\boldsymbol{m}$ vector produces zero torque. The modified law projects $\boldsymbol{m}$ perpendicular to $\hat{\boldsymbol{b}}$ before it gets applied to the magnetic torque coils, which avoids wasting torque coil current.

The second modification is to include integral control terms. The control law computes a vector of error integrals according to the formula

$$\begin{bmatrix} i_r(t) \\ i_p(t) \\ i_y(t) \end{bmatrix} = \frac{1}{\tau_0}\int^t \hat{\boldsymbol{b}}_{\text{des}}(\tau) \times \hat{\boldsymbol{b}}(\tau)\,\mathrm{d}\tau \tag{13}$$

where τ_0 is a controller tuning parameter and τ is a dummy integration variable. The vector cross product is used in the integrand because, for small angular errors, it is equivalent to a sensed vector rotation error.

The roll error integral i_r gets fed back to counteract any steady-state roll disturbance torque, but the pitch and yaw error integrals i_p and i_y are used differently. They prevent the controller from trying to counteract steady-state pitch and yaw torques. It is impossible to fully counteract a general three-axis steady-state disturbance torque with a magnetic control system because of the inability to torque about the instantaneous magnetic field direction. Integral magnetic control can be used to null out the average effects of steady-state torques, but such a strategy excites unwanted periodic oscillations as byproducts.[7] To minimize any such oscillations, the present controller lets the passive aerodynamic stabilization system counteract steady-state pitch and yaw torques. It uses the integral values i_p and i_y to turn off the magnetic feedback actions about the pitch and yaw axes at very low frequencies, so that the magnetic controller will not kick up unnecessary oscillations. This is accomplished by the use of these integrals to define a transformation from spacecraft coordinates to aerodynamic coordinates:

$$A_{\text{aero}}[i_p(t), i_y(t)] = A\left\{\frac{1}{\sqrt{1 + \left[i_p^2(t) + i_y^2(t)\right]/4}}\begin{bmatrix} 0 \\ i_p(t)/2 \\ i_y(t)/2 \\ 1 \end{bmatrix}\right\} \tag{14}$$

where $A\{\boldsymbol{q}\}$ on the right-hand side of Eq. (14) is the 3×3 orthogonal direction cosines matrix associated with the quaternion argument $\boldsymbol{q}$ (Ref. 8). If the controller is working properly, then $A_{\text{aero}}[i_p(t), i_y(t)]$ will settle down to an approximately constant matrix that defines the transformation to a new spacecraft-fixed coordinate system. This coordinate system will get aligned to the local-level coordinate system in steady state; that is, its x axis will point in the velocity direction, and its z axis will point toward nadir. This is called the aerodynamic coordinate system because its x axis points in the velocity direction when the net pitch and yaw disturbance torques about the center of mass are both zero. These disturbance torques are presumed to be the results of aerodynamic imbalance, inertial imbalance, or both.

The rotation in Eq. (14) gets used to transform the measured magnetic field unit direction vector and its time derivative into aerodynamic spacecraft coordinates:

$$\hat{\boldsymbol{b}} = A_{\text{aero}}[i_p, i_y]\hat{\boldsymbol{b}}_{\text{meas}} \tag{15a}$$

$$\dot{\hat{\boldsymbol{b}}} = A_{\text{aero}}[i_p, i_y]\dot{\hat{\boldsymbol{b}}}_{\text{meas}} \tag{15b}$$

where $\hat{\boldsymbol{b}}_{\text{meas}}$ is the measured magnetic field direction vector in the original spacecraft coordinate system. Nonzero time derivatives of i_p and i_y are neglected in computing $\dot{\hat{\boldsymbol{b}}}$ from $\dot{\hat{\boldsymbol{b}}}_{\text{meas}}$ because the true aerodynamic coordinate system should be spacecraft fixed. $A_{\text{aero}}[i_p, i_y]$ has significant time variations only during initial transients of the system, but the integral control terms are not important during the initial transients because of the large values that are normally used for the integral tuning parameter τ_0.

The final form of the control law is

$$\boldsymbol{m} = \left(A_{\text{aero}}^T[i_p, i_y]\left[(I - \hat{\boldsymbol{b}}\hat{\boldsymbol{b}}^T)\left\{\frac{1}{\tau_1}\left(\hat{\boldsymbol{b}}_{\text{des}} + \hat{\boldsymbol{b}} \times \begin{bmatrix} i_r \\ 0 \\ 0 \end{bmatrix}\right) + \dot{\hat{\boldsymbol{b}}}_{\text{des}}\right\} - \dot{\hat{\boldsymbol{b}}}\right]\right)\left(\frac{I_{\text{mavg}}}{\tau_2\|\boldsymbol{b}\|}\right) \tag{16}$$

where $\boldsymbol{m}$ is the magnetic dipole moment vector in the original spacecraft coordinate system, I_{mavg} is the average of the diagonal elements of the spacecraft's moment of inertia matrix, τ_2 is a controller tuning parameter that effectively replaces K_p of Eq. (12), and $\|\boldsymbol{b}\|$ is the norm of the measured magnetic field vector. The leading coordinate transformation matrix on the right-hand side of Eq. (16) transforms the magnetic dipole moment vector from aerodynamic coordinates back into the original spacecraft coordinate system. The projection matrix $(I - \hat{\boldsymbol{b}}\hat{\boldsymbol{b}}^T)$ ensures that $\boldsymbol{m}$ is perpendicular to the magnetic field. The Eq. (16) terms $\hat{\boldsymbol{b}}_{\text{des}}$, $\dot{\hat{\boldsymbol{b}}}_{\text{des}}$, and $\dot{\hat{\boldsymbol{b}}}$ serve the same purposes as the corresponding terms in Eq. (12); they provide proportional and derivative control actions. The Eq. (16) term that involves i_r in a cross product with $\hat{\boldsymbol{b}}$ implements the roll integral control action. The final scalar factor $I_{\text{mavg}}/(\tau_2\|\boldsymbol{b}\|)$ is equivalent to the proportional gain K_p of Eq. (12) multiplied by the prediction time τ_1.

The full control law evaluates Eqs. (14–15b), followed by Eqs. (13) and (16). It contains the tuning parameters τ_0, τ_1, and τ_2. It has been designed so that τ_2 is roughly the time required to

align $\dot{\hat{\boldsymbol{b}}}$ with a perturbation of $\dot{\hat{\boldsymbol{b}}}_{\text{des}}$ that includes a correction term that drives $\hat{\boldsymbol{b}}$ to $\hat{\boldsymbol{b}}_{\text{des}}$. The parameter τ_1 is roughly the time required for this correction term to align $\hat{\boldsymbol{b}}$ with $\hat{\boldsymbol{b}}_{\text{des}}$, and τ_0 is roughly the time required for the integrals in Eq. (13) to reach steady state.

These tuning parameters must be selected properly. No definite theory has been derived to show what values of these parameters will stabilize the system. Such a theory would be difficult to develop due to the time-varying nature of the problem, and a global nonlinear stability proof would require the use of Lyapunov-type techniques. The present work has relied on simulation studies to develop the following ad hoc tuning rules: One should choose $\tau_2 < \tau_1 < \tau_0$. The integral time parameter τ_0 should be on the order of an orbital period or two. The proportional time parameter τ_1 should be tuned to be on the order of the period of the aerodynamic pitch and yaw oscillations or larger. The derivative time parameter τ_2 should equal about one quarter of the proportional time parameter; that is, $\tau_2 = 0.25\tau_1$. If τ_0 is too small, then error averaging assumptions break down, and the integral control actions will not function as envisioned. If τ_1 is too small, then the natural motions of the magnetic field vector as the spacecraft moves along its orbit may cause a parametric excitation of the pitch and yaw dynamics. If τ_1 is too large, however, then roll disturbance torques may not get counteracted fast enough during the parts of the orbit when magnetic roll control is possible.

C. Control Saturation and Integrator Antiwindup Logic

The control law must deal wisely with actuator saturation. Suppose that the maximum torque coil dipole moment for each axis is m_{sat}. Control saturation logic scales back the magnitude of the $\boldsymbol{m}$ vector from Eq. (16) if the magnitude of any of its elements exceeds m_{sat}:

$$\boldsymbol{m}_{\text{act}} = \begin{cases} \boldsymbol{m} & \text{if} \quad \|\boldsymbol{m}\|_\infty < m_{\text{sat}} \\ \left(\dfrac{\boldsymbol{m} m_{\text{sat}}}{\|\boldsymbol{m}\|_\infty}\right) & \text{if} \quad m_{\text{sat}} \le \|\boldsymbol{m}\|_\infty \end{cases} \tag{17}$$

where $\boldsymbol{m}_{\text{act}}$ is the actual dipole moment vector that gets sent to the torque coils and where $\| \ \|_\infty$ is the usual infinity norm. This saturation logic leaves the direction of $\boldsymbol{m}$ unchanged, thereby preserving the compasslike actions that form the rationale behind the control law in Eq. (16).

Antiwindup logic is needed to avoid stability problems that can occur if the integrators build up to excessive values during times of actuator saturation. The antiwindup logic turns off the integration in Eq. (13) whenever $m_{\text{sat}} \le \|\boldsymbol{m}\|_\infty$, that is, whenever the lower condition in Eq. (17) applies. During these times the integrals $i_r(t)$, $i_p(t)$, and $i_y(t)$ remain constant.

D. Pulse-Width-Modulated Implementation and Power Usage

The magnetic dipole moment vector $\boldsymbol{m}$ will normally be implemented with pulse-width modulation. Pulse-width modulation is easy to implement electronically; it requires only a timer and switching circuitry. Under this scheme, the j-axis coil is turned on at the value $\text{sign}[(\boldsymbol{m}_{\text{act}})_j] m_{\max}$ for $T_{\text{sat}}|(\boldsymbol{m}_{\text{act}})_j|/m_{\text{sat}}$ seconds during each pulse cycle. In these formulas, $m_{\max}$ is the full-on coil strength and T_{sat} is the maximum coil on time during a pulse cycle. The duration of each pulse cycle is $T_p = m_{\max} T_{\text{sat}}/m_{\text{sat}}$. The condition $T_{\text{sat}} < T_p$ is enforced to allow the torque coils to be turned off for a brief interval during each pulse cycle, so that the magnetometer can be read without cross talk from the coils. This control scheme results in an average dipole strength of $(\boldsymbol{m}_{\text{act}})_j$ during the pulse cycle. If T_p is chosen to be small compared to the system's open-loop and closed-loop time constants, then the closed-loop response under pulse-width-modulated control will be nearly identical to what the response would have been if the magnetic dipole moment had been commanded to the value $\boldsymbol{m}_{\text{act}}(t)$ on a continuous basis.

The system's average power usage is an important metric of its practicality. The power usage of the torque coils under pulse-width-modulated control can be computed based on the coils' properties. The average power during a single pulse cycle is

$$P_{\text{avg}} = \frac{m_{\max} \|\boldsymbol{m}_{\text{act}}\|_1 R_c}{N_c^2 A_c^2} \tag{18}$$

where N_c is the number of turns in a coil, A_c is the area of a coil, R_c is the resistance of a coil, and $\|\boldsymbol{m}_{\text{act}}\|_1$ is the 1-norm of $\boldsymbol{m}_{\text{act}}$, which equals the sum of the absolute values of the elements of $\boldsymbol{m}_{\text{act}}$. Note that a pulse-width-modulated system uses more power than a system that allows continuous variation of the applied magnetic dipole moment.

IV. Simulation of Closed-Loop System Response

Several aspects of this design need to be evaluated before an actual system is flown. First, one needs to verify the global stability of the combined aerodynamic/magnetic stabilization system. This includes the ability to stabilize from an initial tumbling situation and the ability to damp out oscillations enough to achieve reasonably small peak pointing errors in steady state. Controller tuning can affect stability, and reasonable tuning values need to be determined. Various disturbances act on the system. These include mass unbalance, aerodynamic twist of the feathers, rotation of the atmosphere with the Earth (which disturbs yaw for inclined orbits), and magnetometer errors. Parametric and classical resonance have been found to be important for such systems,[2] and their effects on the present system must be investigated. Another important issue concerns the maximum useful altitude for such a system. Above some altitude, the air density becomes too low, and this system does not provide adequate stability.

These issues have been investigated by simulation of the system's response. The simulation uses Euler's equation for a rigid body and the quaternion kinematics equation to model the closed-loop attitude dynamics. The forcing terms in Euler's equation include the torques due to aerodynamics, the magnetic torque coils, the gravity-gradient effect, solar radiation pressure, and radiation pressure from the Earth's albedo.

The simulation's atmospheric model is based on the 1976 U.S. Standard Atmosphere.[6] The altitude input to this model is computed with the World Geodetic Survey-84 ellipsoid model of the Earth. The atmosphere is assumed to rotate with the Earth, which causes aerodynamic yaw disturbances in inclined orbits. The model does not include atmospheric winds. The model includes an ad hoc diurnal density bulge that elevates the density by a factor of four at 1400 hrs local time at the same latitude as the sun. This bulge drops off as a two-dimensional Gaussian with a standard deviation of 52 deg in longitude and 29 deg in latitude. Some simulations include an additional ad hoc density variation along the orbit to simulate the effects of a sharper diurnal bulge. This extra variation gives the density time history sufficient amplitude at high multiples of the orbital frequency to test whether the system can undergo parametric resonance, as in Ref. 2. This extra variation is needed for two reasons: first, because the ad hoc diurnal bulge yields a smoother density time history than Ref. 2 finds with the physics-based Jacchia–Roberts atmospheric model, and second, because the present paper's stiffer aerodynamic design can experience a principal parametric resonance only if it gets forced at relatively high multiples of the orbital frequency.

The radiation torque models are similar to those used in Refs. 1 and 2. They use a panel model of the spacecraft and tabulate torque divided by incident radiation pressure as a function of the spacecraft-relative azimuth and elevation angles of the incident radiation. These tabulated models account for shading and include specular reflection, diffuse reflection, and absorption. The online torque calculations consider eclipse of the sun by the Earth. If the sun is not eclipsed, then the incident pressure and direction are calculated as functions of spacecraft position and attitude and are used in conjunction with bilinear interpolation of the table to calculate the torque. An Earth albedo model is included. It assumes that the albedo radiation is concentrated as specular reflection from the appropriate point on the Earth with a reflected radiation pressure that is equal to 25% of the incident radiation pressure, except that the albedo pressure decays to zero over a short portion of the orbit when the spacecraft goes into eclipse. Otherwise, the albedo radiation pressure torque calculation works like the solar radiation pressure torque calculation, which uses the table, the incident direction, and the incident pressure to calculate the torque.

The simulation's orbital dynamics model includes Keplerian dynamics and secular $J2$ effects. It does not include the effects of drag, solar radiation pressure, periodic $J2$ terms, higher-order gravity potential terms, sun–moon effects, or tidal effects. This is reasonable because the goal of the simulation is to observe the medium-term attitude behavior on the scale of several orbits. This simulation is known to be accurate to within about 3 km or better for such cases, which is sufficiently accurate for purposes of the present study.

Magnetometer errors are important because they affect the system through feedback in the control law of Eqs. (13–16). The simulation includes the following magnetometer errors: scale-factor, bias, relative alignment, absolute alignment, and random noise. Relative misalignment refers to nonorthogonality of the sensitive axes. The scale-factor, bias, and relative-alignment errors are sized to be what one would expect after an attitude-independent calibration of a typical flight magnetometer. The random-noise errors and the absolute misalignment between the magnetometer axes and the spacecraft axes are chosen to be representative of existing flight units.

V. Simulation Results

A number of simulations have been run to test the performance of this system. Consider the results for a typical case. This case corresponds to the baseline shuttlecock spacecraft design of Sec. II.D. Its moment of inertia matrix is $I_m = \text{diag}(0.00211, 0.00952, 0.00952)\ \text{kg}\cdot\text{m}^2$, and its center of mass is balanced perfectly on the nominal roll axis. Its feathers have zero collective twist, so that there is no aerodynamic roll torque. The spacecraft orbits in a circular orbit with an altitude of 400 km and zero inclination. The magnetic controller design parameters have been set at $\tau_0 = 6000$ s, $\tau_1 = 800$ s, and $\tau_2 = 200$ s. The saturation dipole moment for each torque coil axis is $m_{\text{sat}} = 0.1\ \text{A}\cdot\text{m}^2$. The magnetometer has a total misalignment of 0.35 deg, nonorthogonality errors ranging from 0.05 to 0.41 deg, scale factor errors ranging from 0.02 to 0.38%, per-axis biases ranging from 166 to 904 nT, a $1-\sigma$ noise of 8.7 nT per axis, and a digitization granularity of 12.2 nT. Although the simulation does not use a pulse-width-modulated implementation, the following parameters apply to the power that would be used by an equivalent pulse-width-modulated implementation that allowed the coils to be turned on for a maximum of 80% of a pulse cycle, that is, for $T_{\text{sat}} = 0.8T_p$: $m_{\max} = 0.125\ \text{A}\cdot\text{m}^2$, $A_c = 0.0064\ \text{m}^2$, $N_c = 18$, and $R_c = 0.22\ \Omega$.

Figure 5 shows the results for this case. Figure 5a shows the roll, pitch, and yaw angle time histories, and Fig. 5b shows the roll, pitch, and yaw rate time histories. The initial conditions for this case correspond to a tumbling situation. The initial roll, pitch, and yaw rates are, respectively, −3, −2, and 2.5 deg/s. The initial attitude angles are −13 deg in roll, 101 deg in pitch, and −26 deg in yaw.

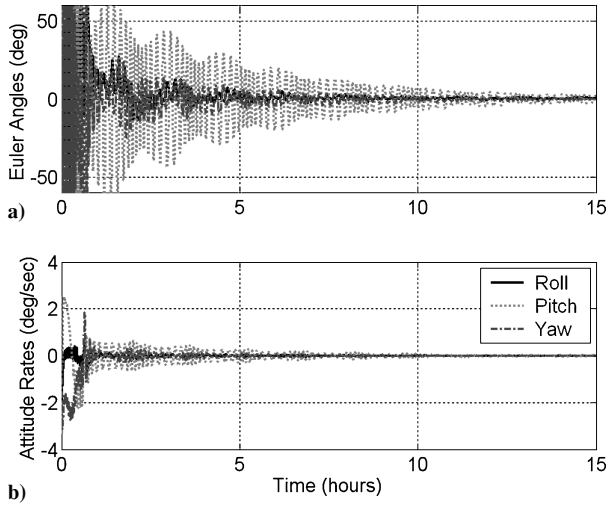

Fig. 5 Closed-loop attitude and rate time histories starting from a tumbling condition, a typical case.

Figure 5 demonstrates that the aerodynamic pitch–yaw stabilization system and the magnetic feedback control system work together to globally stabilize the attitude. The system stops tumbling within 1 h, as indicated by Fig. 5b. The pointing errors are less than 7 deg per axis after 10 h, and after 15 h they are about 2 deg. The maximum power demand of this system is 0.21 W during the initial transient, and the mean power use in steady state is 0.0024 W.

Figure 5 shows a problem that one encounters when using magnetic stabilization at low orbital inclinations: The magnetic field vector stays nearly aligned with the nominal pitch axis. Its average and peak angular deflections from the nominal pitch axis are 19 and 27 deg, respectively. This near alignment makes the system less able to add damping about the pitch axis, which is why the pitch oscillations take the longest to die out in Fig. 5. Fortunately, the Earth's magnetic field never remains exactly aligned with the nominal pitch axis in any orbit. Otherwise, the magnetic control system would not be able to damp out pitch oscillations. Three factors that preclude steady alignment of the field with the nominal pitch axis are 1) the cant of the Earth's main dipole field with respect to the Earth's rotation axis, 2) the rotation of the Earth, and 3) higher-order field irregularities.

A challenging stabilization case uses the same nominal shuttlecock aerodynamic design as the preceding case, but it includes collective twist of the feathers and center-of-mass imbalance. The former effect causes a roll aerodynamic disturbance torque, and the latter effect causes pitch and yaw aerodynamic disturbance torques. The collective twist is 20 deg for each feather, and the mass imbalance of the main spacecraft bus is 0.025 m along the y and z spacecraft axes. These represent worst-case values.

Additional challenges for the next case are posed by its orbital conditions. The orbit is circular with an altitude of 457 km and an inclination of 87 deg. This altitude is significantly higher than for the preceding case. The pitch and yaw undamped natural frequencies equal 4.7 and 5 times the orbital frequency, respectively. These conditions are right for parametric and classical resonance of the yaw mode, as discussed in Ref. 2. The likelihood of parametric resonance is further enhanced by artificially increasing one of the periodic terms in the atmospheric density time history. The term whose frequency equals 10 times the orbital frequency has its zero-to-peak density variation artificially increased to equal 30% of the mean density along the orbit. In the notation of Ref. 2, this corresponds to a μ_{10} value of 0.15. This large value of μ_{10} will give rise to a strong instability in the principal region of parametric resonance if the active pitch–yaw damping system is ineffective. Yet another challenge is the high orbital inclination of this case. The rotation of the atmosphere with the Earth and the high inclination combine to produce periodic forcing of the yaw mode at the orbital frequency.

The other parameters of this case are defined as follows: The controller tuning values, the magnetometer measurement error model, and the torque coil properties are the same as in the preceding case. The moment-of-inertia matrix is slightly different for this case due to the changed location of the main bus center of mass:

$$I_m = \begin{bmatrix} 0.00214 & -0.00021 & -0.00026 \\ -0.00021 & 0.00898 & 0.00009 \\ -0.00026 & 0.00009 & 0.00898 \end{bmatrix} \text{kg}\cdot\text{m}^2 \qquad (19)$$

Results for this challenging case are presented in Fig. 6, which shows the roll, pitch, and yaw angle time histories. The spacecraft achieves full three-axis stabilization when it starts from the tumbling initial condition of 3, 3.5, and −4 deg/s for the roll, pitch, and yaw rates, respectively. Tumbling ceases in less than 1 h. The system reaches steady state in about 10 h. After this, its maximum pointing errors are 22 deg in roll, 13 deg in pitch, and 27 deg in yaw. The large roll errors are primarily due to the aerodynamic twist. The large yaw errors are due to the parametric and classical resonance and to the atmospheric rotation effects. Pitch and yaw also display steady-state biases of 7.8 and −3.5 deg, respectively, roughly consistent with the mass imbalance and the form of the pitch and yaw integral action. The roll bias is less than 1 deg, which indicates that the

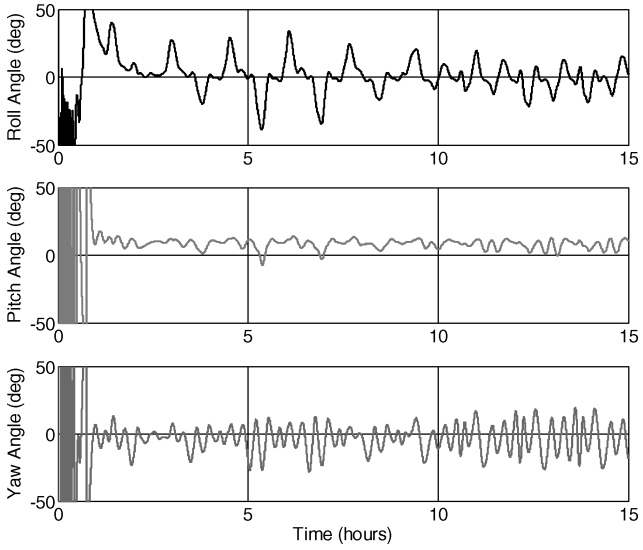

Fig. 6 Closed-loop attitude time histories in an inclined orbit with aerodynamic twist, mass imbalance, and parametric and classical resonance.

control law's roll integrator works properly. Although not shown in any figure, the pitch error integral $i_p(t)$ settles to within 0.4 deg of the true value, which the system's mass imbalance causes to be 7.5 deg, and the yaw error integral $i_y(t)$ settles to within 2.5 deg of its true value, which is 8.6 deg. Thus, the pitch and yaw integrators function properly. The average steady-state power usage for this case is 0.0018 W. This is lower than for the case associated with Fig. 5, which is surprising given the larger disturbances of the present case.

A number of other simulation cases have been run, and the following discussion highlights important insights that have been learned from these cases. Successful stabilization has been achieved at altitudes up 497 km with the same controller tuning values. The tuning values $\tau_0 = 6000$ s, $\tau_1 = 400$ s, and $\tau_2 = 100$ s produce instability at this altitude, despite the fact that they yield stability at lower altitudes. This underscores the point made at the end of Sec. III.B: There are lower bounds on τ_0, τ_1, and τ_2 that must be respected to achieve stability. These bounds increase as altitude increases because of the increased period of the open-loop pitch and yaw oscillations.

Steady-state oscillation amplitudes in response to disturbances tend to increase with increases in the orbit's altitude. This is true because an increased altitude decreases the open-loop bandwidth of the pitch–yaw aerodynamic stabilization system; the aerodynamic stiffness is proportional to atmospheric density,[2] and density decreases as altitude increases. As an example, two cases have been run at two different circular orbit altitudes, 400 and 497 km. Both cases have an inclination of 87 deg, and both use the same tuning parameters as given for the Fig. 5 case. The only disturbances are those due to the rotating atmosphere, solar and albedo radiation pressure, and magnetometer measurement errors. There is no mass imbalance, no collective twist of the feathers, no open-loop parametric resonance, and no classical resonance. At 400-km altitude, the peak steady-state pointing errors are 11.7 deg in roll, 2.0 deg in pitch, and 7.1 deg in yaw, but at 497 km, the peak pointing errors are 21.7 deg in roll, 7.4 deg in pitch, and 35.6 deg in yaw. Thus, an increased altitude increases the pointing errors of the closed-loop system.

The effects of the controller tunings on the steady-state pointing errors are more complex. The simulation results suggest that there is an optimal tuning. If τ_0, τ_1, and τ_2 are too large or too small, then steady-state pointing errors increase from the optimum. The nominal values of $\tau_0 = 6000$ s, $\tau_1 = 800$ s, and $\tau_2 = 200$ s are the best tuning values for the example spacecraft that have been found to date. Consider what happens when the faster tuning values $\tau_0 = 6000$ s, $\tau_1 = 400$ s, and $\tau_2 = 100$ s are used in an orbit with an altitude of 400 km and an inclination of 87 deg. These values nearly double the peak steady-state roll error, and they increase the peak yaw excursions by more than 50%. At 497-km altitude and 87-deg inclination, the slower tuning values $\tau_0 = 12{,}000$ s, $\tau_1 = 2000$ s, and $\tau_2 = 500$ s cause the peak steady-state roll and pitch pointing errors to grow by 35 and 82%, respectively. A possible explanation for the existence of optimal tuning values is that damping becomes too low when the values are too small, but bandwidth becomes too low when the values are too large.

These results can be summarized as follows. The shuttlecock design with magnetic feedback can be three-axis stabilized up to 500-km altitude if the controller tuning gains are chosen properly. Atmospheric rotation, magnetometer measurement errors, spacecraft mass imbalance, collective aerodynamic twist of the feathers, solar radiation pressure torque, and parametric resonance from periodic atmospheric density variations can cause steady-state pitch and yaw biases and steady-state oscillations on all three axes. The nominal system considered in this study can be designed to be globally stable below 500 km and to have maximum per-axis pointing errors of 36 deg or less. These pointing errors decrease with decreases in the altitude, the orbital inclination, the mass imbalance, the aerodynamic twist of the feathers, and the magnetometer measurement errors.

VI. Conclusions

A new type of three-axis attitude stabilization system has been designed for a low-Earth-orbiting nadir-pointing spacecraft. It uses passive aerodynamic drag torques to stabilize the pitch and yaw axis and active magnetic control to stabilize the roll axis and to provide damping in pitch and yaw. The aerodynamic configuration resembles a badminton shuttlecock and provides adequate stabilization at altitudes below 500 km. The magnetic control law uses proportional–integral–derivative output feedback of magnetometer readings to magnetic torque coils. It points the spacecraft, so that the measured magnetic field direction in spacecraft coordinates matches the desired value. The desired value is computed by the use of the orbital position, an Earth magnetic field model, and the desired attitude.

The performance of this system has been investigated via simulation. The aerodynamic simulation uses a free-molecular flow model of the atmosphere's interactions with the spacecraft's surfaces. The following issues have been studied: the system's global stability properties, the influence of controller tuning, the possibility of parametric resonance due to the diurnal bulge in the Earth's atmosphere, and the disturbance effects of manufacturing tolerance errors, solar radiation pressure torques, gravity-gradient torques, and the rotation of the Earth's atmosphere.

The resulting system is lightweight and performs reasonably well for a 1-kg-class satellite. Its feathers, feather deployment system, magnetometer, and magnetic torque coils can be built and operated within a 0.115-kg mass budget and a 0.06-W average power budget, and its feedback calculations can be carried out with only a small fraction of the command and control processor's capacity. Its worst-case steady-state pointing error is 25 deg/axis for altitudes below 450 km. This error bound will decrease if the orbital inclination, the altitude, or the allowable manufacturing errors decrease.

Acknowledgments

Valuable assistance on this project was provided by Iain D. Boyd of the University of Michigan and by David R. Skillman of NASA Goddard Space Flight Center. I. D. Boyd introduced the author to Ref. 5 and provided comments about the aerodynamic model. D. R. Skillman pointed the author to the work in Refs. 1–3 and provided a copy of Ref. 3.

References

[1]Kumar, R. R., Mazanek, D. D., and Heck, M. L., "Simulation and Shuttle Hitchhiker Validation of Passive Satellite Aerostabilization," *Journal of Spacecraft and Rockets*, Vol. 32, No. 5, 1995, pp. 806–811.

[2]Kumar, R. R., Mazanek, D. D., and Heck, M. L., "Parametric and Classical Resonance in Passive Satellite Aerostabilization," *Journal of Spacecraft and Rockets*, Vol. 33, No. 2, 1996, pp. 228–234.

[3]Heck, M. L., "Final Report: Analysis of PAMS Flight Data," Report 97-6, Analytical Mechanics Associates, Inc., Hampton, VA, March 1997.

[4]Guelman, M., Flohr, I., Ortenberg, F., Shachar, M., Shiryaev, A., Volfovsky, A., and Waler, R., "The Israeli Microsatellite TechSat for Scientific and Technological Research: Development and In-Orbit Testing," *Acta Astronautica*, Vol. 46, Nos. 2–6, 2000, pp. 397–404.

[5]Gombosi, T. I., *Gaskinetic Theory*, Cambridge Univ. Press, Cambridge, England, U.K., 1994, pp. 227–255.

[6]United States Committee on Extension to the Standard Atmosphere, *U.S. Standard Atmosphere, 1976*, National Oceanic and Atmospheric Administration, Washington, DC, 1976, pp. 50–97.

[7]Psiaki, M. L., "Magnetic Torquer Attitude Control via Asymptotic Periodic Linear Quadratic Regulation," *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 2, 2001, pp. 386–394.

[8]Wertz, J. R. (ed.), *Spacecraft Attitude Determination and Control*, D. Reidel, Boston, 1978, p. 414.