Adaptable Method of Managing Jets and Aerosurfaces for Aerospace Vehicle Control

Joseph A. Paradiso*
Charles Stark Draper Laboratory, Inc., Cambridge, Massachusetts 02139

An actuator selection procedure is presented that uses linear programming to optimally specify bounded aerosurface deflections and jet firings in response to differential torque and/or force commands. This method creates a highly adaptable interface to vehicle control logic by automatically providing intrinsic actuator decoupling, dynamic response to actuator reconfiguration, dynamic upper bound and objective specification, and the capability of coordinating hybrid operation with dissimilar actuators. The objective function minimized by the linear programming algorithm is adapted to realize several goals, i.e., discourage large aerosurface deflections, encourage the use of certain aerosurfaces (speedbrake, body flap) as a function of vehicle state, minimize drag, contribute to translational control, and adjust the balance between jet firings and aerosurface activity during hybrid operation. A vehicle model adapted from Space Shuttle aerodynamic data is employed in simulation examples that drive the actuator selection with a six-axis vehicle controller tracking a scheduled re-entry trajectory.

Introduction

THE complex missions and demanding environment considered for tomorrow's generation of aircraft and aerospace vehicles will impose increasingly formidable challenges on candidate control schemes. These vehicles will require control laws that can exploit the full potential of all available actuators in order to adapt quickly to changing vehicle characteristics, while maintaining stringent constraints on vehicle state. Prime examples can be found in proposed hypersonic vehicles, which are intended to perform as aircraft from takeoff through at least the initial portion of ascent. At extreme altitudes, aerosurfaces become ineffective; hence, these vehicles must be controlled as spacecraft, via a reaction control system (RCS) and propulsive thrust vector control. The sequence reverses upon descent, where the RCS is initially needed for vehicle stabilization, with a gradual transition to aerodynamic control after the aerosurfaces gain sufficient authority. Throughout the atmosphere flight, propulsive and thermal considerations impose strict constraints on vehicle attitude and aerosurface deflection.

Control systems partially addressing this challenge have been developed to manage re-entry of the Space Shuttle.[1] To handle the transition from RCS to aerodynamic control as dynamic pressure increases, the current Shuttle autopilot uses several different control strategies that are sequentially applied at different points during the descent.

Managing each group of actuators with independent control logic can result in reduced vehicle controllability and efficiency. Because aerospace vehicles need to combine the actions of various types of actuators during both ascent and descent in order to cope with variations in dynamic pressure and air-breathing engine operating characteristics, they will require a highly coordinated actuator management scheme. An adaptive hybrid control strategy is needed that is capable of extracting maximum performance from each actuator family (in solo performance or concerted operation) and optimally reconfiguring during evolution of the vehicle environment or after hardware failures. Such reliability will be flight-critical, as even a transient degradation in control at high Mach number could result in loss of the vehicle.

Improvements in available onboard computing capacity have enabled aircraft and spacecraft to employ control schemes of increasing complexity. An actuator management system based on real-time execution of linear programming has produced a highly adaptable fuel-optimal jet selection,[2] which has already been successfully flight tested[3] onboard the Shuttle Orbiter. Additional studies have revised and extended these concepts to incorporate control moment gyroscopes (CMGs) into the selection process.[4] By dynamically adjusting objective factors, upper bounds, and failure flags associated with each set of actuators, linear programming can adaptively determine efficient and effective policies of actuator usage. Since all available actuators are considered together in a common "pool," this technique has the ability to select and blend the action of various types of effectors (i.e., jets, aerosurfaces, propulsion), resulting in true "hybrid control."

Previous aircraft control efforts[5] have employed a pseudoinverse solution to linearly map desired body torques into aerosurface commands. Such methods can provide control laws with intrinsic longitudinal/lateral actuator decoupling, yet the conventional pseudoinverse calculation lacks the capabilities provided by linear programming to impose hard constraints on actuator usage and establish actuator preference via an objective function. Incorporating features such as these in pseudoinverse and conventional schemes could imply considerable adaption of the control laws, which may become less feasible after actuator failures and reconfiguration, leading to potentially degraded performance. Linear programming retains the benefit of intrinsic actuator decoupling, while providing the control logic with the ability to dynamically specify the preferred actuator behavior and limit actuator displacement.

An additional benefit of this approach is the potential of coordinating both translational and rotational vehicle response simply by extending the order of actuator activity vectors (i.e., measures of vehicle response to specific actuator motion) to also account for translational degrees of freedom. In this fashion small corrections to the flight path can be accommodated by allowing the selection to specify actuator lift and drag while maintaining full rotational control. Because this scheme can account for all degrees of freedom simultaneously, it is intrinsically able to compensate for cou-

Received July 17, 1989; presented as Paper 89-3429 at the AIAA Guidance, Navigation, and Control Conference, Boston, MA, Aug. 14–16, 1989; revision received Nov. 21, 1989.

*Member of Technical Staff, Sensors, Signal, and Data Processing Directorate. Member AIAA.

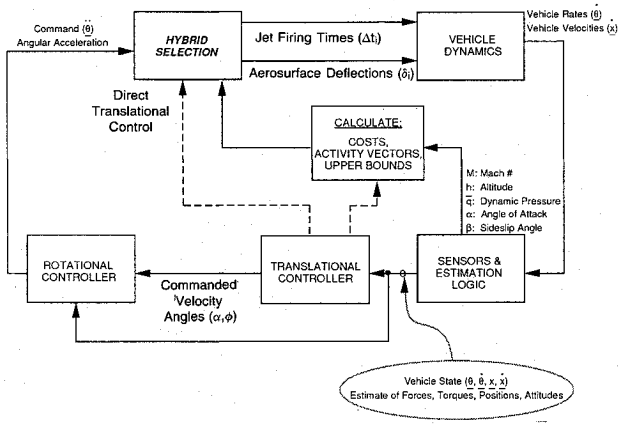

Fig. 1 Hybrid control scheme as applied to aerospace vehicle.

pled translational and rotational response to actuator deflection. Decoupled control of vehicle translation and rotation has been investigated in experimental aircraft (i.e., the AFTI/F-16; Ref. 6) and promises widespread potential application.

Figure 1 shows a diagram depicting the means by which a hybrid selection procedure can be integrated into an aerospace plane flight control package. A collection of sensors, along with appropriate estimation algorithms and software, is assumed to provide a dynamic measurement of the vehicle state (attitudes, rates position, velocity) and environment (forces, torques, aerosurface/jet authorities, dynamic pressure, etc). These quantities are used to continually update parameters for the linear selection, i.e., aerosurface activity vectors (estimate of instantaneous torque/force control authorities), costs (objective penalizations per actuator), and upper bounds (maximum allowed deflections per control step). In order to compensate for aerodynamic effects or changing vehicle mass properties, activity vectors representing jet acceleration may also be periodically updated as a function of vehicle state.

An estimate of vehicle position and velocity is compared with a set of desired values in a translational controller, which generates a command velocity attitude (angle of attack α and bank ϕ; sideslip β is generally held at zero), which will correct the net force on the vehicle. The translational control logic is also able to input a translational force change command directly to the hybrid selection (leftmost dotted line in Fig. 1), allowing the actuators themselves to directly deliver the requested force difference. One must bear in mind, however, that the aerosurfaces and jets are only capable of restricted translational authority due to the limited aerosurface area and constraints on available jet thrust and fuel. Provided that there are sufficient independent aerosurfaces available to maintain simultaneous rotational control, this option may prove useful for translational trimming, as demonstrated in the simulations of Ref. 7.

Linear Programming Formation

The "hybrid selection" package incorporated in Fig. 1 executes a linear program to determine the optimal mix of bound aerosurface deflections and jet firings that yield a commanded vehicle response. An estimate of the instantaneous rotational and translational authorities of all actuators (termed "activity vectors") is scanned during the selection process. Each actuator possesses at least one associated objective coefficient, upper bound, and failure flag that determine its desirability, authority limit, and availability, respectively. The linear programming problem solved by the hybrid selection may be summarized as follows:

Minimize

$$Z = \sum_{j=1}^{N} c_j |x_j| \tag{1}$$

Subject to

$$\sum_{j=1}^{N} \boldsymbol{A}_j x_j = \Delta \boldsymbol{R} \tag{2a}$$

$$-U_j^- \le x_j \le U_j^+ \tag{2b}$$

where

N = number of actuators available to the system
c_j = cost factor associated with actuator j
$U_j^{\pm}$ = upper/lower bounds associated with actuator j
$\boldsymbol{A}_j$ = activity vector representing authority of actuator j
x_j = decision variable denoting action of actuator j
$\Delta \boldsymbol{R}$ = requested vehicle acceleration change

Equation (2a) states an equality constraint. It is a vector equation representing an underdetermined system of M scalar equations (M = number of controlled axes, i.e., dimension of $\boldsymbol{A}_j$ and $\Delta \boldsymbol{R}$) in N unknowns. Equation (2b) is an inequality constraint expressing independent upper and lower bounds on the allowed range of the decision variables x_j. By employing the "upper bounding simplex algorithm,"[8] the limits of Eq. (2b) can be considered without augmenting the order of the problem stated in Eq. (2a).

Equation (1) is the linear objective function that is minimized in the solution to the linear program; it essentially defines a weighted 1 norm in the space of decision variables x_j. The solution values of x_j denote the selected amounts of corresponding actuator action (i.e., change in aerosurface deflection). Limits on actuator usage may be imposed independently in either direction by clamping positive and negative decision values by their corresponding bounds ($U_j^{\pm}$). The activity vectors in this framework ($\boldsymbol{A}_j$) denote the instantaneous acceleration change estimated to be produced by each actuator per unit decision value x_j. The vector $\Delta \boldsymbol{R}$ is the acceleration change command input from the vehicle controller.

The algorithm used to solve the linear program has been adapted from the revised simplex method applied in Ref. 9. In particular, modifications have been introduced to make simplex consult different activity vectors for opposite-sign aerosurface deflection (accounting for potential difference in aerosurface response moving into and out of the airstream) and to avoid violating tank constraints on jet firings. The simplex implementation has also been fashioned to enable dynamic control of up to six axes; i.e., any combination of rotational and translational coordinates may be placed under (or removed from) direct actuator control at any time by setting (resetting) an appropriate Boolean flag. The details of the specific simplex implementation used in this study are presented in Ref. 7.

Activity vectors, decision variables, and upper bounds are defined as follows for aerosurfaces (subscript i) and jets (subscript j).

Activity vectors:

$$\boldsymbol{A}_i^{\pm} = \begin{bmatrix} [I]^{-1}\, \mathrm{d}\boldsymbol{\tau}_i^{\pm} \\ 1/M\, \mathrm{d}\boldsymbol{F}_i^{\pm} \end{bmatrix} \qquad \boldsymbol{A}_j = \begin{bmatrix} [I]^{-1} \boldsymbol{r}_j \times \boldsymbol{T}_j \\ 1/M\, \boldsymbol{T}_j \end{bmatrix} \tag{3a}$$

where the equation on the right-hand side is for jets, and the equation on the left-hand side is for aerosurfaces. The upper three activity vector components (above the dashed line) describe rotational authority, whereas the lower three components (below the dashed line) describe translational authority.

Decision variables:

$x_i = \Delta \delta_i$ = aerosurface deflection change,

$x_j = D_j$ = jet duty cycle (3b)

Upper bounds:

$$U_i^{\pm} = \min \begin{cases} \pm L_i & U_{(j)}^{+} = 1.0 \\ \delta^{\pm}_{\text{stop}(i)} - \delta_i & U_{(j)}^{-} = 0 \\ \dot{\delta}^{\pm}_{\max(i)} \Delta t_c \end{cases} \tag{3c}$$

where

$[I]$ = estimate of vehicle inertia matrix
M = estimate of vehicle mass
$d\tau_i^{\pm}$ = torque authority of aerosurface i in $\pm$ direction
$d\boldsymbol{F}_i^{\pm}$ = force authority of aerosurface i in $\pm$ direction
δ_i = current deflection of aerosurface i
$\delta_{\text{stop}(i)}$ = stop location for aerosurface i
$\dot{\delta}_{\max(i)}$ = maximum slew rate for aerosurface i
Δt_c = control update period
$\boldsymbol{r}_j$ = position of jet relative to the vehicle center of mass
$\boldsymbol{T}_j$ = thrust of jet j

Aerosurface activity vectors specify the change in airframe rotational and translational acceleration expected per unit deflection of each corresponding aerosurface (linearized about the current aerosurface position for each sign of deflection, i.e., $\pm$). The aerosurface decision variables specified by the linear program are the changes in deflection angles. These are clamped by the upper bounds expressed at left in Eq. (3c); the most restrictive bound currently evaluated in the bracket is applied.

The topmost item in the bracket of Eq. (3c) (i.e., L_i) is a generic clamp on allowed deflection change per control step. The middle expression represents the angle between the current aerosurface deflection and the maximum "stop limit" (in the appropriate $\pm$ direction). This limits the absolute deflection angle and prevents the linear selection from prescribing a deflection change that places an aerosurface beyond its allowed range. The variable $\delta^{\pm}_{\text{stop}(i)}$ may be varied dynamically, allowing the restriction on aerosurface deflections to evolve during various flight regimes. The bottom expression in Eq. (3c) represents the maximum deflection possible per control time step (Δt_c). This limits the participation of various aerosurfaces in the solution in order to account for the different slew rates attainable by each actuator. These factors may be specified differently for opposite-sense deflection, as indicated by the $\pm$ superscript.

Jets are defined as continuous torque actuators under the selection framework. The jet accelerations (angular and translational for up to 6-DOF control) are used as activity vectors, as summarized at right in Eq. (3a), and in correspondence with the conventions pursued in Refs. 2 and 4. The jet decision variables (x_j) are now defined to be jet duty cycles (as opposed to jet firing times, which were used in the previous efforts). They range from 0 to 1 and define the fraction of maximum jet acceleration needed to realize the input command. These continuous duty cycles are realized as closely as possible by discrete jet firings in the simulated vehicle environment. The ratio of average jet *on* times to *off* times is made proportional over the control update interval to the corresponding duty cycles (discretized, however, by the minimum allowed jet firing times). By setting upper bounds to unity for jet decision variables (and lower bounds to zero through the intrinsic "feasibility" constraint), simplex will solve directly for jet duty cycles in response to an acceleration-change input command.

The effects of any currently firing jets are not considered in the estimate of vehicle acceleration used in computing the commanded acceleration change [$\Delta \boldsymbol{R}$ in Eq. (2)]. This causes the jet commands to be absolute, i.e., all jets are initialized to be off at the start of each selection, and absolute duty cycles are specified when jets are required. Vacuum jet accelerations are assumed in the activity vectors of Eq. (3a); effects arising from perturbations due to aerodynamic jet interaction are examined in Ref. 7.

The linear program may select other actuators by defining suitable activity vectors (here representing expected acceleration-change effect), upper bounds (limiting actuator participation to reflect constraints on actuator range, bandwidth, and authority), objective factors (indicating the desirability of using each actuator), and decision variables (through which the actuators are commanded). A means of extending this framework to specify main engine throttle level and thrust vector rotation for management of an ascent vehicle has been proposed in Ref. 7.

Objective Formulation

The objective function minimized by simplex [Eq. (1)] is a sum of weighted cost contributions. Terms are included to discourage needless jet activity, penalize aerosurface deflection angle, and avoid maximum deflection limits (i.e., stops). The objective calculation for each actuator may be summarized as follows:

$$c_j = K_{\text{jet}}(j) \tag{4a}$$

where the activity vector $\boldsymbol{j}$ corresponds to an RCS jet,

$$c_{j,s} = K_0(j) + K_A F_{\text{angle}}(j,s) + K_S G_{\text{stops}}(j,s) + K_T V_{\text{translation}}(j,s) + K_Q Q_{\text{specific}}(j,s) \tag{4b}$$

where the activity vector $\boldsymbol{j}$ corresponds to an aerosurface, and s is the sign of deflection, i.e. $\pm$.

The objective penalization of RCS jets is given by a single term, K_{jet}. This factor is different for various sets of jets (i.e., the use of forward jets is penalized more heavily, since they can appreciably perturb entry aerodynamics) and is altitude-dependent (jets are made more expensive as the vehicle descends and are eventually prohibited altogether at low altitude, see Ref. 7). The K_{jet} factors are generally significantly higher than average aerosurface costs, in order to discourage jet firings except where absolutely necessary.

The cost calculation for dynamic actuators (such as aerosurfaces) is, however, more complicated and includes terms from several sources. The leading term, K_0, is a bias that dictates the general desirability of using a particular actuator. If K_0 is relatively large, the actuator will be avoided in a solution (where possible), with its participation increasing as K_0 drops. The F_{angle} and G_{stops} functions act to penalize deflection. In general, one would like to keep aerosurfaces near their trim positions (in the absence of other considerations). This is encouraged for small and moderate deflections through the F_{angle} function, which adds an amplitude into the objective penalizing simplex solutions that increase deflection angles:

$$F_{\text{angle}}(j,s) = \begin{cases} |\delta_j| & \text{if rotation } s \text{ increases } |\delta_j| \\ 0 & \text{otherwise} \ldots \end{cases} \tag{5}$$

Deflection increments that increase the magnitude of net deflection angle $|\delta_j|$ are assigned a cost contribution in direct proportion to the current value of $|\delta_j|$. Deflections that decrease $|\delta_j|$ are given no cost contributions via F_{angle}. Rotations that increase the deflection angle thus become linearly more expensive as the angle grows. Solutions involving the activity vector and decision variables that bring $|\delta_j|$ back to zero accordingly become increasingly favored as $|\delta_j|$ rises.

If an actuator is pinned against a hard stop, a degree of freedom is essentially lost to the selection algorithm (the actuator can then only be moved in one direction, i.e., off the stop). In addition, thermal and hinge-moment constraints may create regions near the extremes of aerosurface deflections that should be avoided whenever possible. Although the upper bounds of Eq. (3c) may be imposed to absolutely prevent actuator motion past stop boundaries, an objective function that increases rapidly as an actuator nears its limit

could slow or inhibit actuator motion before maximum deflection is reached. Advancement of aerosurfaces to large deflections will thus be allowed, but strongly discouraged, and not selected except when absolutely necessary to maintain vehicle control.

The G_{stops} cost contribution signals such a warning to the selection procedure as an actuator nears its limit. In contrast to the linear form of F_{angle}, G_{stops} contributes a nearly insignificant amount to the objective if the aerosurface is removed from its stop [allowing the other terms in Eq. (4b) to act unimpeded], but increases rapidly after the aerosurface has approached to within a preset distance from the stop location. The form of G_{stops} chosen to be applied here can be expressed as

$$G_{\text{stops}}(j,s) = \begin{cases} \Lambda(\delta_j) & \text{if rotation } s \text{ moves actuator toward stop} \\ 0 & \text{otherwise} \ldots \end{cases} \tag{6}$$

where

$$\Lambda(\delta_j) = \tan\left[\frac{\pi}{2}\left\{(1-\zeta)\left[\frac{\delta_j}{\delta_{\text{stop}(j)}}\right]+\zeta\right\}\right] - \tan\left(\frac{\pi}{2}\zeta\right)$$

$$\zeta = \text{steepness paremeter}; \qquad 0 < \zeta < 1$$

The function Λ has a small value for low δ_j; however, as $\delta_j/\delta_{\text{stop}}$ approaches unity, Λ diverges asymptotically to infinity. One may control the break point at which Λ diverges by adjusting the ζ parameter in Eq. (6). In general, ζ is chosen near 0.93, placing the break point in aerosurface deflection roughly at $\delta_{\text{break}} \approx 0.75\,\delta_{\text{stop}}$.

If the rotation s brings an aerosurface toward a stop, the objective contribution will be proportional to Λ. No such contribution will be added to the objective coefficient if the aerosurface is defined to have unlimited freedom, or if rotation s will remove it from a stop. If an actuator has neared its stop, the function Λ will contribute appreciably, and solutions that rotate the actuator away from the stop are heavily favored in contrast to those that move it closer. The form of Λ in Eq. (6) may be simplified (one can use several divergent functions); it was set up in its present realization to facilitate modifications during testing. Both functions F_{angle} and G_{stops} attempt to minimize deflection angles, but the steep G_{stops} contribution works primarily at large δ_j, whereas the function F_{angle} has a significant effect at smaller δ_j.

The inclusion of translational effect into the cost function was defined by the $V_{\text{translation}}$ term of Eq. (4b). This amplitude can be defined to aid in longitudinal control, which can be stated as

$$V_{\text{translation}}(i,\pm) = K_x \Delta D(A_i^{\pm})_x - K_z \Delta L(A_i^{\pm})_Z \tag{7}$$

where ΔD is the desired change in drag force; ΔL is the desired change in lift force; and $A_i^{\pm}$ is the activity vector for aerosurface i in $\pm$ direction (x and z components, defined in stability axes).

Equation (7) assigns a cost contribution to each aerosurface activity vector in proportion to its authority in drag and lift. Deflection is encouraged in a direction to produce the desired effect, with an urgency proportional to the magnitude of the requested change. Longitudinal control can be accommodated by setting ΔL and ΔD to the changes in lift and drag requested by a translational controller. A common bias term is added to keep the $V_{\text{translation}}$ factors positive. A preference for minimizing aerosurface drag can be injected into the objective, per Eq. (7), by setting K_z to zero and ΔD to -1. Equation (7) could also provide lateral control by adding a third y term in a similar fashion.

The objective contribution of Eq. (7) encourages aerosurface deflection to produce a gross translational effect. Equation (7) is not a hard constraint; thus, it only expresses a "desire" for a translational force change. The redundancy in the actuator system is exploited to achieve a specified translational response, whereas the commands specified in the equality constraint of Eq. (2b) are met exactly. Since the separate aerosurfaces have smaller authority than the full airframe, this technique is primarily useful for small translational trimming or achieving a generic effect (such as minimizing actuator drag).

The Q_{specific} term in Eq. (4b) represents an objective amplitude that encourages a desired behavior of particular aerosurfaces; this term can be exploited to achieve dynamic actuator desaturation and track scheduled aerosurface deflection profiles, as detailed in the following section.

Vehicle Model and Implementation

To investigate the performance of the actuator selection procedure using established data describing a hypersonic vehicle, simulations conducted during this study have adopted a model based on the standard Space Shuttle aerosurface and jet configurations as defined for re-entry. The vehicle adopted in these tests is assumed to possess seven controllable aeroactuators. Two elevons, a body flap, a rudder, and a speedbrake are incorporated as conventionally defined,[10] with parameters summarized in Ref. 7. In addition, two canards were added to the model as a means of increasing the alternatives available to the actuator selection. Although the canards are not needed for conventional three-axis attitude control, tests that investigate failure cases and attempt simultaneous actuator control of rotational and translational vehicle states benefit from the extra degrees of freedom. A diagram depicting the location of the seven aerosurfaces is given in Fig. 2. Positive elevon, canard, and body flap deflections are defined as moving down into the airflow at positive angle of attack.

The inner and outer panels of the left and right elevons are assumed to always deflect equally (as is the convention during Shuttle entry), forming a single effective elevon on each side of the vehicle. Recent studies[11] indicate that differential deflection of inboard and outboard elevon panels can provide a means of controlling vehicle yaw at high angle of attack (where the rudder is ineffective), reducing the need for jet firings. The hybrid selection is entirely capable of specifying this; indeed, differential deflection would be performed automatically to provide yaw, if inboard and outboard panels were separately available to the linear program. Because the vehicle model used here does not enable independent inner/outer elevon control, this capability can not be demonstrated in these results. However, a similar effect is possible by differentially deflecting the elevons and canards; the linear program is, in fact, seen[7] to exploit this possibility for gleaning additional yaw authority.

Since the Space Shuttle lacks canards, this control contribution is approximated by scaling the reaction to an equivalent

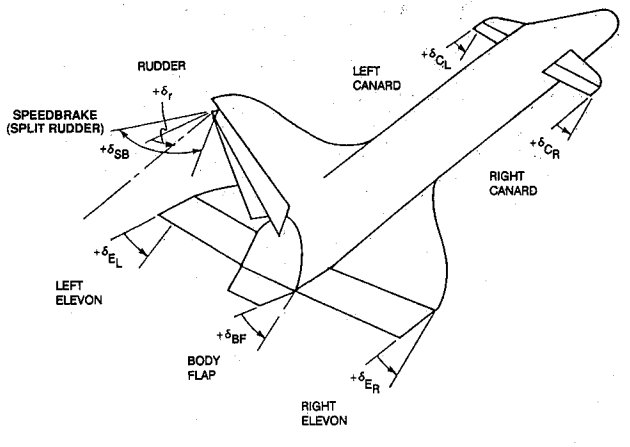

Fig. 2 Modeled aerosurface locations and sign conventions.

deflection of the corresponding elevon by -1 in pitch (since these canards are assumed to be placed considerably forward of the vehicle c.g.) and by 0.1 in roll, yaw, and translational forces (primarily due to their smaller aerosurface area). Canards slew at 20 deg/s (as do the elevons), but their range of deflection is more restricted (± 10 deg for canards vs -35 deg $\rightarrow +20$ deg for elevons). Admittedly, the analogy between canards and elevons is a crude one; this elementary formulation is intended only for demonstrating the performance of the hybrid selection and control procedure with an additional set of aeroactuators.

The speedbrake on the Shuttle vehicle is realized by a split rudder; both surfaces open symmetrically under speedbrake deflection. Because of their correlated operation, the maximum allowed rudder deflection can depend on the current speedbrake angle (and vice versa). These constraints are dynamically accommodated by the linear program, as described in Ref. 7.

Coefficients describing the forces and torques exerted on the airframe as a function of vehicle attitude, Mach number, and aerosurface deflection were constructed from an extensive collection of aerodynamic data describing Shuttle re-entry.[12] These coefficients were reduced to a network of data points that describes the aerodynamic action on the vehicle at various attitudes, airspeeds, and actuator deflections within the Shuttle operational envelope. An efficient multidimensional interpolation procedure is then invoked by the hybrid controller to consult this table of sampled data in real time and estimate the aerodynamic forces, torques, and aerosurface authorities at the current vehicle state, which are, in turn, used to update the vehicle control logic and form actuator activity vectors.

The vehicle state and actuator authorities input to the control logic are generally taken directly from the output of the environment software. No model of sensor hardware or state estimator performance is inserted into the data flow. However, some error is naturally introduced through inherent aerodynamic nonlinearity; i.e., linear predictions from tabulated data at the current attitude can be somewhat different several time steps later after the vehicle rotates. A quick investigation into the effect of estimation uncertainty is presented through a set of examples in Ref. 7 that examine the vehicle and controller response to random and systematic modeling errors. Adaptive state estimation and dynamic identification algorithms for an aerospace vehicle are currently under development.[13]

All selections assume both jets and aerosurfaces to be available (provided the vehicle is not below the 45,000-ft jet cutoff altitude). Methods of limiting the nonlinear perturbations from wide aerosurface deflections are described in Ref. 7.

The ability to set independent objective coefficients for each aerosurface (and sign of deflection) has been exploited to tailor the action of certain aerosurfaces (i.e., body flap and speedbrake; see Fig. 2) to specific applications. A major function of the body flap is to reduce the combined elevon (elevator) deflection during Shuttle re-entry. The objective function has been adapted here [through the Q_{specific} term of Eq. (4b)] to automatically extend the body flap in order to relieve the elevon and canard load. Body flap deflections leading to reduced elevon/canard loads are assigned a negative cost value. This encourages appropriate body flap deflection to be selected (thus yielding smaller elevon/canard angles) until its excursion becomes appreciable (causing the stops and deflection costs to contribute significantly, thus canceling the negative body flap cost), or the elevons and canards return to trim.

The speedbrake has very limited authority across most of the regime studied in these tests and (especially with the presence of canards) is not needed to complete commands. In order to adequately exhibit its use, however, a series of tests have been performed[7] that dynamically assign the speedbrake a high negative cost at certain times during the re-entry (again through Q_{specific}) to encourage its deflection and schedule its deployment in accordance with standard Shuttle practice.[14]

A two-level control hierarchy was developed to drive the linear selection in simulation examples. At the highest level, a translational controller uses estimates of longitudinal and lateral position errors to produce angle-of-attack and bank commands, which, at the lower level, are realized by a three-axis rotational controller. Although both translational and rotational controllers are based on gain-scheduled proportional-integral-derivative (PID) compensators, the actuator selection process may be easily adapted for other control schemes; these simple controllers were constructed only to demonstrate the hybrid selection procedure through re-entry simulations. Details of the control logic are presented in Ref. 7. The vehicle assumed in these studies is considered to act as a rigid body. Flexible dynamics are not currently applied in either the control design or simulation dynamics.

A control update rate of 1.5 Hz was found to provide adequate stability and limit perturbation due to nonlinearities in simulation examples. However, the presence of increased modeling uncertainty, estimation effects, and a less benign vehicle environment may necessitate a higher repetition rate. This should be possible to achieve; a linear program (running revised simplex without bounded variables) has already been cycled at up to 12.5 Hz as an experimental jet selection onboard the Shuttle Orbiter.[3] Since upper bounds may be imposed on decision variables without excessive additional computation, the linear programming algorithm used here should be able to execute at a similar rate (the upper-bounded simplex algorithm[8] accommodates bounds by checking the response of the solution to candidate activity vectors; there is no need to augment the order of the equality constraint through artificial variables). Faster execution should be easily attainable in newer, more powerful processors, and some portions of the simplex algorithm (i.e., scanning of activity vectors for invitation into the solution) may be effectively parallelized.

Simulation Example

A series of examples[7] have explored the performance of the hybrid control approach. The simulation presented here investigates a portion of vehicle re-entry, tracking a Shuttle-derived trajectory from 170,000 ft at Mach 12 through approximately 20,000 ft at Mach 0.5. At the higher altitudes (and larger α values), both jets and aerosurfaces are required for vehicle control. At the lower altitudes (and smaller α values), aerosurfaces are capable of maintaining control without jet assistance, and the vehicle can be managed conventionally as an unpowered aircraft. A major advantage of the linear programming approach, as demonstrated by the re-entry tests, is its ability to readily adapt to the changing aerodynamic conditions encountered across the entry trajectory; i.e., a single control scheme can essentially manage the vehicle through several aerodynamic regimes. In addition to following the scheduled longitudinal profile, the vehicle is commanded to momentarily displace laterally by 10,000 ft during the descent.

The jet/aerosurface cost balance has been adjusted to avoid excessive aerosurface chatter and respond with a mainly jet-based solution when the aerosurface authority is limited. The speedbrake is assigned a large cost, effectively discouraging its deployment. Positive canard deflection is prohibited during the initial portion of this test (i.e., through the first 400 s) by setting its upper bound [U^+, Eq. (2b)] to zero. After this interval, the canards are allowed to deflect up to their full 10-deg range. Thermal and aerodynamic considerations may impose such dynamic constraints on aerosurface operation in an actual hypersonic vehicle as various flight regimes are encountered. Although effects such as minimizing actuator drag projection can be achieved by ma-

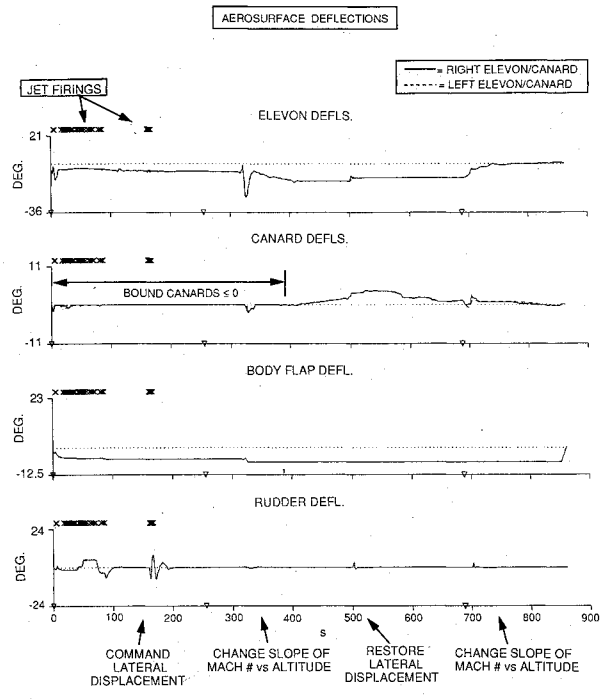

Fig. 3 Simulation example: aerosurface deflections.

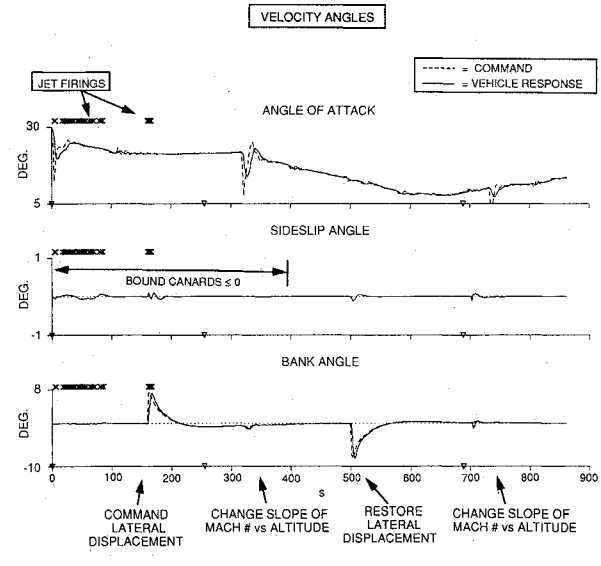

Fig. 4 Simulation example: velocity angles.

nipulating the $V_{\text{translation}}$ contribution [see Eq. (7)], it is set to zero in this example; only deflection minimization and stops avoidance terms are retained in Eq. (4b) (except the body flap, which has a Q_{specific} dynamically defined to unload the elevons and canards).

Aerosurface deflections are given in Fig. 3, velocity angles in Fig. 4, and translational states in Fig. 5. The ×'s at the left of these plots (above the curves) indicate that jets are required at high α (where the rudder is shadowed by the airframe); hybrid operation continues until α drops below roughly 23 deg, and the rudder gains yaw authority. Jets are also needed to aid in producing the +7-deg bank at 150 s, forming the side velocity for the 10,000-ft lateral translation. Jets are not required for the negative bank at 500 s (which restores lateral position); since α is much lower here (i.e., slightly under 10 deg), the rudder has considerable authority and can stabilize yaw unaided.

All of these jet firings were purely lateral (no pitch compo-

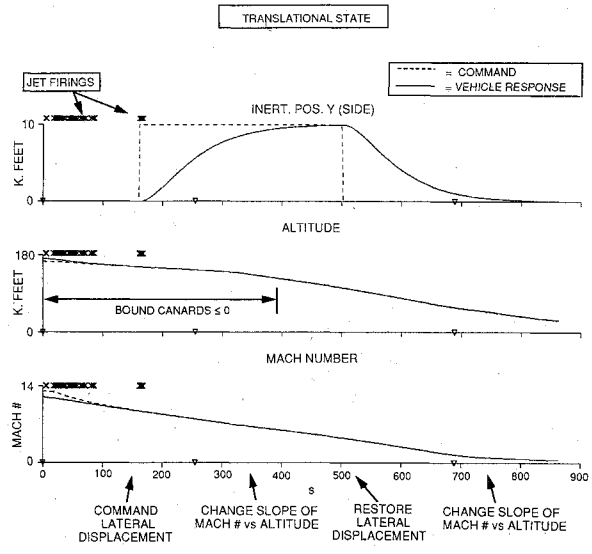

Fig. 5 Simulation example: translational states.

nent is present); aft side-firing jets were chosen exclusively. Although other jets were available for selection, the much less expensive aerosurfaces had ample authority in pitch and roll; thus, jets were only applied to stabilize yaw. The aerosurfaces possess limited means of attaining yaw torque independent of rudder deflection, i.e., opposing canard/elevon scissoring can produce some decoupled yaw effect (although this mode is inhibited somewhat at the beginning of this test by the prohibition imposed on positive canard deflection). The tradeoff between excessive aerosurface deflection to produce a small amount of needed torque vs the introduction of jet firings is managed effectively by adjusting jet/aerosurface cost balance, as illustrated through a series of examples presented in Ref. 7. The aft jets chosen in this example also produce considerable roll acceleration, which must be compensated by additional aerosurface deflection. Forward jets (which can produce yaw with less roll effect) were also available to the selection at this altitude, but were not chosen because of their higher cost (assigned due to the larger potential for aerodynamic perturbation). This example does not consider the effect of aerodynamic flow on jet firings. Initial studies[7] have indicated that such effects are most significant along the roll axis (especially for the aft jets fired here) and can be readily countered by the aerosurfaces as a disturbance torque.

Jet firings are quantized to a minimum duration of 80 ms in this simulation to discretely realize the duty cycles output from the linear program. By introducing hysteresis through deadband techniques, such as a yaw phase-plane controller, the collection of small firings exhibited at high α could be replaced with fewer discrete firings of larger magnitude.

The vehicle is seen to precisely follow the commanded α profile in Fig. 4 (the dashed curve representing the commanded α state is almost completely overdrawn by the solid curve representing the vehicle response). Angle of attack is seen to track a relatively smooth profile, except for commanded impulses at the start of the test (where α is modulated to null the initial altitude and Mach errors), at roughly 350 s midway through the test (where α is pulsed after the vehicle passes Mach 7 to quicken the rate of vertical descent in correspondence with a discrete increase in slope of the commanded altitude vs Mach number profile), and at roughly 800 s near the end of the test (these impulses are due to disturbances encountered near Mach 1, together with a reduction in slope of the commanded altitude vs Mach number profile). These α impulses were not preprogrammed; they result directly from the longitudinal control logic responding to translational state errors. The two bank maneuvers evident in Fig. 4 were commanded by the lateral controller to attain

and remove the velocity needed for the 10,000-ft side excursion. Sideslip disturbances were effectively rejected throughout the simulation.

Figure 5 indicates that the vehicle translational state followed the commanded trajectories (dotted curves). The lateral impulse was successfully achieved, and the desired longitudinal profile was closely tracked after nulling initial state errors.

As seen in Fig. 3, aerosurfaces systematically deflect throughout the test to offset the gradual evolution in aerodynamic environment with changing α and Mach number (except the more rapid response needed during hybrid maneuvers and at points where the commanded translational trajectory changes). The deflection plots of Fig. 3 are scaled to accommodate the maximum allowed aerosurface angles. Significant rudder deflection is needed to stabilize yaw at high α (although much of the control burden was carried here by jet firings) and to produce the initial bank needed for the lateral displacement. Later in the test (at lower α), the required yaw control is achieved with a much smaller deflection.

The body flap is seen to initially deflect upward in an effort to unload the elevons and canards, as driven by Q_{specific}. Note the manner in which the body flap deflection returns to zero, in correspondence with the reduction in combined elevon and canard angles at the conclusion of the test. Because the speedbrake was not deployed, its deflection history is not presented.

The effect of the bound on positive canard displacement is evident in Fig. 3. Simulations run with the canards allotted their full deflection range attempt to balance initial vehicle pitch torques between both elevons and canards (i.e., elevons are deflected negative and canards are deflected positive) because of the relative parity between their pitch authorities and cost magnitudes. Since the bound initially imposed on the canards limits their participation, the elevons are deflected additionally to deliver the needed torque. This reliance on elevon activity can be seen explicitly in the actuator response to the commanded α-impulse at 320 s. A significant negative elevon swing was needed to realize the vehicle pitch command; in tests run without the canard bound,[7] positive canard participation greatly reduced the needed elevon excursion. Since negative canard deflection was not precluded throughout this example, negative canard swings were occasionally selected to aid in achieving the commanded α impulses and trimming jet disturbances. After the positive canard bound is restored to its 10-deg maximum at 400 s, the canards promptly begin drifting into positive deflection to better assist in balancing pitch torques. Significant elevon and canard deflection is reached by 550 s into the test to counter an increase in pitch torque at lower α and higher dynamic pressure. Aerodynamic conditions at the conclusion of the test enable the aerosurfaces to balance torques without straying far from zero deflection.

Conclusion

This study has demonstrated that linear programming holds the potential to perform as a highly adaptable actuator management procedure. In addition to the intrinsic control axis decoupling provided by the linear program, dynamic constraints on actuator usage can be enforced by imposing upper bounds, and desired actuator behavior (i.e., actuator preference, deflection minimization, drag avoidance, desaturation, and scheduling) can be effectively encouraged through appropriate definition of an objective function. By specifying relevant activity vectors, bounds, objectives, and decision variables, the linear program can be made capable of choosing an optimal mix of actuators to answer acceleration-change commands and thus direct hybrid multiactuator vehicle operation. These concepts have been demonstrated by adapting a linear programming algorithm to effectively select jet firings and specify aerosurface deflections for re-entry control of a simulated aerospace vehicle.

Acknowledgments

The author is highly grateful to Ed Bergmann, Phil Hattis, and Darryl Sargent of the Charles Stark Draper Laboratory for helpful advice and fruitful discussions. This effort was supported under Charles Stark Draper Laboratory Independent Research and Development Task 236.

References

[1]"Operational Flight Level C, Functional Subsystem Software Requirements (FSSR); Guidance, Navigation, and Control, Part C, Flight Control Entry—GRTLS," Rockwell/NASA, Rept. STS 83-0007, 1983.

[2]Bergmann, E. V., Croopnick, S. R., Turkovich, J. J., and Work, C. C., "An Advanced Spacecraft Autopilot Concept," *Journal of Guidance and Control*, Vol. 2, No. 3, 1979, pp. 161–168.

[3]Bergmann, E. V., "OEX Advanced Autopilot Flight Test Overview," Society of Automotive Engineers Aerospace Control and Guidance Systems Committee, Meeting 61, Monterey, CA, March 1988. [The OEX Autopilot was flight-tested on Space Shuttle missions STS 51G (June 1985) and STS 61B (Nov. 1985).]

[4]Paradiso, J. A., "A Highly Adaptable Steering/Selection Procedure for Combined CMG/RCS Spacecraft Control," AAS 86-036, Feb. 1986.

[5]Raza, S. J., and Silverthorn, J. T., "Use of the Pseudo-Inverse for Design of a Reconfigurable Flight Control System," AIAA Paper 85-1900, Aug. 1985.

[6]Speyer, J. L., White, J. E., Douglas, R., and Hull, D. G., "MIMO Controller Design for Longitudinal Decoupled Aircraft Motion," *Proceedings of the AIAA Guidance and Control Conference*, AIAA, New York, 1983.

[7]Paradiso, J. A., "Application of Linear Programming to Coordinated Management of Jets and Aerosurfaces for Aerospace Vehicle Control," Charles Stark Draper Lab., Cambridge, MA, Rept. CSDL-R-2065, Nov. 1988.

[8]Bradley, S. P., Hax, A. C., and Magnanti, T. L., *Applied Mathematical Programming*, Addison-Wesley, Reading, MA, 1977.

[9]Paradiso, J. A., "A Highly Adaptable Steering/Selection Procedure for Combined CMG/RCS Spacecraft Control, Detailed Report," Charles Stark Draper Lab., Cambridge, MA, Rept. CSDL-R-1835, March 1986.

[10]*Operational Aerodynamic Data Book*, Vols. 1 and 3, Rockwell Space Div., Downey, CA, Rept. STS85-0118, Sept. 1985.

[11]Brown, L. W., and Montgomery, R. C., "Space Shuttle Separate-Surface Control-System Study," NASA T-P 2340, N84-28801, July 1984.

[12]Silver, L. W., "ESIM Model Book for the C. S. Draper Laboratory Statement Level Simulator (Revision 4)," Charles Stark Draper Lab., Cambridge, MA, Rept. R-776, June 1980 (the particular data used in the aerodynamic calculation was taken from SLS REL 9).

[13]Chamitoff, G., "Intelligent Adaptive Flight Control for Hypersonic Vehicles," PhD Dissertaiton, Dept. of Aeronautics and Astronautics, Massachusetts Inst. of Technology, Cambridge, MA, Jan. 1991.

[14]Findlay, J. T., Kelley, G. M., Heck, M. L., and McConnell, J. G., "Summary of Shuttle Data Processing and Aerodynamic Performance Comparison for the First Eleven (11) Flights; Final Report," NASA CR-172440, N87-10884, Sept. 1987.