Sensitivity Analysis of Linear Programming and Quadratic Programming Algorithms for Control Allocation

Susan A. Frost[1]
NASA Ames Research Center, Moffett Field, CA 94035

Marc Bodson[2]
University of Utah, Salt Lake City, UT 84112

and

Diana M. Acosta[3]
NASA Ames Research Center, Moffett Field, CA 94035

The next generation (NextGen) transport aircraft configurations being investigated as part of the NASA Aeronautics Subsonic Fixed Wing Project have more control surfaces, or control effectors, than existing transport aircraft configurations. Conventional flight control is achieved through two symmetric elevators, two antisymmetric ailerons, and a rudder. The five control surfaces, reduced to three command variables, produce moments along the three main axes of the aircraft and enable the pilot to control the attitude of the aircraft. Next generation aircraft will have additional redundant control effectors to control the three moments, creating a situation where the aircraft is over-actuated and where a simple relationship no longer exists between the required control surface deflections and the desired moments. NextGen flight controllers will incorporate control allocation algorithms to determine the optimal effector commands to attain the desired moments, taking into account the effector limits. Approaches to solving the problem using linear programming and quadratic programming algorithms have been proposed and tested. It is of great interest to understand their relative advantages and disadvantages and how design parameters may affect their properties. In this paper, we investigate the sensitivity of the effector commands with respect to the desired moments and show on some examples the sensitivity of the solutions provided by the linear programming and quadratic programming methods.

[1] Research Engineer, Intelligent Systems Division, M/S 269-1, and AIAA Member.
[2] Professor, Electrical & Computer Engineering Department, 50 S Central Campus Dr Rm 3280, and AIAA Senior Member.
[3] Research Engineer, Intelligent Systems Division, M/S 269-1, and AIAA Member.

I. Introduction

Concern for the environmental and operational impacts of civil aviation is prompting significant changes in civil aviation. The NASA Aeronautics Subsonic Fixed Wing Project, together with industry, universities, and other government organizations, is researching advanced technologies and exploring novel civil transport configurations to achieve the desired reduction in noise, emissions, fuel burn, and field length for Next Generation (NextGen) aircraft[1-2]. Potential NextGen civil transport aircraft include the Cruise Efficient Short Take-Off and Landing (CESTOL) aircraft and the Hybrid Wing-Body (HWB) aircraft. See ref. 2 for detailed information on these new NextGen architectures and their control challenges.

The advanced configurations and associated technologies for next generation aircraft elevate the complexity of the vehicle and its operation. Flight control will be particularly challenging, in part because of the added demands on the control allocation schemes. Unlike conventional control allocation schemes that control three body-fixed rates primarily with three control variables, NextGen control allocation schemes will need to control the three body-fixed rates with a variety of redundant and multi-objective control surfaces[3]. The control allocation of over-actuated aircraft has been formulated as a constrained optimization problem by some researchers[4]. For some NextGen configurations, a real-time solution to the control allocation problem is desirable to enable the system to run on the aircraft during flight. This has prompted the search for numerical optimization methods that have good convergence properties and acceptable computational requirements. Other NextGen aircraft configurations may require a predetermined mapping between desired accelerations and control surface deflections, which could be optimally determined using these same constrained optimization methods.

Several methods to solve the control allocation problem have been evaluated, including direct allocation, linear programming, quadratic programming, weighted pseudo-inverse, cascaded generalized inverse, and a mixed optimization approach[4,5]. Control allocation research has also extended the control solution to include coupling or interaction effects between control effectors, creating a nonlinear optimization problem that can often be transformed into a linear problem[6,7]. While the interaction effects will be important to study for control allocation in next generation aircraft, this paper will focus on solutions that assume a linear relationship between the effectors and the moments they generate.

Two optimization methods that have been shown to be effective and implementable in real-time are the linear programming (LP) and the quadratic programming (QP) methods[4,5]. Previous studies, however, have not addressed the issue of the sensitivity of the solution to the data. In particular, the sensitivity of the solution with respect to the desired moments determines how rapidly the control surfaces have to move if the moments are to match the desired moments and the moments change rapidly. The sensitivity of two control allocation algorithms, one using an LP method and the other using a QP method, are evaluated in this paper by examining their application to an example of an over-actuated aircraft.

II. Control Allocation in Flight Control Systems

Control allocation is the stage in the stability and control augmentation process that is responsible for achieving a commanded moment vector by allocating commands to individual control effectors. Flight control algorithms for next generation aircraft will control q body-fixed angular accelerations or moments of the aircraft with n control effectors, where $n > q$.

In its simplest practical form, the control allocation scheme will consider the rate and magnitude limits of the control effectors. The control allocation problem can be posed in state space form as follows: find the $m \times 1$ control vector, u_{cmd}, such that

$$Bu_{cmd} = \dot{x}_{cmd} \tag{1}$$

subject to

$$u_{\min} \le u_{cmd} \le u_{\max} \text{ and } \left|\dot{u}_{cmd}\right| \le \dot{u}_{\max} \tag{2}$$

where B is the $q \times m$ linear control effectiveness matrix, $\dot{x}_{cmd}$ is the $q \times 1$ commanded angular acceleration vector, $u_{\min}$ and $u_{\max}$ are the vectors of the control effectors' minimum and maximum magnitude limits, respectively, and $\dot{u}_{\max}$ is the vector of control effectors' maximum rate limits. In this paper, we set aside the rate limits and focus on the accelerations or moments that are achievable, without considering the previous deflections that may have been applied.

The control allocation algorithm should optimally allocate the available control effectors to solve eqs. (1-2), while limiting unachievable commands. The attainable moment set (AMS) for a given operating point of an aircraft is the set of moments, or accelerations, that can be achieved by the control effectors without violating their saturation limits. If the solution to eqs. (1-2) is not unique, additional objectives, such as control minimization, can be added to the optimization problem to yield a preferred solution.

A. Mixed Optimization Approach to Control Allocation

The mixed optimization approach, as proposed by Bodson[5], combines the error minimization and the control minimization into one problem by minimizing the cost function:

$$J = \left\|Bu_{cmd} - \dot{x}_{cmd}\right\|_1 + \varepsilon\left\|u_{cmd} - u_{pref}\right\|_1 \tag{3}$$

subject to

$$u_{\min} \le u_{cmd} \le u_{\max} \tag{4}$$

where $u_{\min}$ and $u_{\max}$ are the vectors of the control effectors' minimum and maximum magnitude limits, $0 \le \varepsilon < 1$ is a control minimization weight, and u_{pref} is a preferred control vector. If the value of the parameter ε is small, priority will be given to error minimization over control minimization, as desired. Note that the optimization criterion is based on the l_1 norm, which is the sum of the absolute values of the components of the vector.

The mixed optimization approach expands upon work performed by Buffington[9], who suggested that control allocation could be solved using a sequential optimization approach in which the error minimization problem is first solved followed by solution of the control minimization problem. The specific choice of l_1 norm enabled Buffington to solve the control allocation problem with linear programming techniques, providing guaranteed convergence to a solution in an acceptable period of time.

The mixed optimization approach, as implemented by Bodson[5], combined the two optimization steps into one, and was shown to be the most effective LP optimization method in an evaluation of different approaches used in simulations of a C-17 aircraft model and a tailless aircraft model. The algorithm was based on the revised simplex method[10] with additional refinements as described in detail in ref. 5, which enabled computations to be performed extremely fast, and well within the capabilities of modern computers. This paper uses the same code for the LP optimization method in the comparison studies that follow.

B. Active Set Methods for Solving Quadratic Control Allocation

Quadratic programming or Least-Squares (LS) methods make use of the l_2 norm and have been used extensively in the control allocation problem[4,11-13]. Härkegård[11] suggested the use of active set methods[8] for quadratic solutions to the control allocation problem. Active set methods can be shown to find an optimal solution in a finite number of steps[11]. A comparison of constrained quadratic programming techniques for control allocation found that the weighted least-squares (WLS) form of the active set method performed favorably for control allocation of over-actuated aircraft that resemble next generation aircraft configurations[12]. The active set algorithm is to quadratic programming what the simplex algorithm is to linear programming. The steps of the algorithm are very comparable, and the solutions can be computed very fast with both methods.

The general problem formulation for the quadratic solution to the mixed optimization control allocation problem is to minimize the cost function:

$$J = \left\| Bu_{cmd} - \dot{x}_{cmd} \right\|_2^2 + \varepsilon^2 \left\| u_{cmd} - u_{pref} \right\|_2^2 \tag{5}$$

subject to

$$u_{\min} \le u_{cmd} \le u_{\max} \tag{6}$$

where $u_{\min}$ and $u_{\max}$ are the vectors of the control effectors' minimum and maximum magnitude limits, $0 \le \varepsilon < 1$ is a control minimization weight, and u_{pref} is a preferred control vector. Note that the optimization criterion in eq. (5) is comparable to the criterion in eq. (3), with the l_2 norm being used instead of l_1 norm. This leaves open the question of whether the choice of norm makes any difference in the solutions obtained. The point of this paper is to show that it does. The WLS algorithm[11,13] was used as the QP method for analysis in this paper. See ref. 13 for the source code and documentation on the WLS algorithm.

III. Computational Results

Computations were performed on a aircraft model based on Lockheed Martin's ICE (innovative control effectors) tailless aircraft[9,15]. The aircraft has 11 control effectors: left elevon, right elevon, pitch flaps, left all-moving tip, right all-moving tip, pitch thrust vectoring, yaw thrust vectoring, left spoiler slots, right spoiler slots, left outboard leading-edge flaps, and right outboard leading-edge flaps. The control effectiveness matrix, or *B* matrix, corresponding to an output vector composed of pitch, roll, and yaw rates for a flight condition at Mach 0.4 and 15,000 ft altitude is given in table 1. The position limits of the control effectors are given in table 2. Note that the limits of the spoiler slot deflectors were lowered from 60 degrees to 10 degrees in ref. 15 to reduce nonlinear interactions between the spoiler slot deflectors and the elevons. The same limits were used here, although the nonlinear effects were not part of the evaluation. In the computations, the value of ε used was $\varepsilon = 10^{-3}$ and $u_{pref} \equiv [0]_{11x1}$.

In order to quantify the sensitivity of the algorithms to the data, a sensitivity metric was computed to assess how much the solution changes in a neighborhood of a given input, $\dot{x}_{cmd}$. The sensitivity metric compares the computed effector deflections for a given desired moment vector and the computed deflections for the same desired moment vector with a small constant vector added to it. The sensitivity for an acceleration vector $\dot{x}_{cmd}$ is given by

$$\text{sens}(\dot{x}_{cmd}) = \frac{\left\| u_{cmd} - u_{delta} \right\|_2}{\left\| \Delta \right\|_2} \tag{7}$$

where u_{cmd} is the vector of control surface deflections computed by the algorithm for $\dot{x}_{cmd}$ and u_{delta} is the vector of control surface deflections for $\dot{x}_{cmd} + \Delta$, with Δ a small acceleration vector.

Figures 1-4 compare the results of simulations for the ICE with a pure pitch acceleration command given as the desired acceleration, *e.g.,* the roll and yaw acceleration components were identically zero. The simulation input had the pitch acceleration start at zero and step up to 1.1 times the maximum attainable pitch acceleration of 249.23 deg/s^2 for the ICE model. There were 110 points in the interval. For this simulation, Δ always had the pitch acceleration equal to one and the roll and yaw accelerations equal to zero. This value of Δ was added to the desired acceleration to test the sensitivity at each point. Figure 1 compares the sensitivity of the two algorithms, demonstrating that the sensitivity of the LP algorithm is slightly greater than the sensitivity of the QP algorithm for this simulation.

Figure 2 shows the commanded control surface deflections for the ICE model using the LP and QP algorithms. Comparison of the plots in fig. 2 reveals that the allocation of the control surfaces was more distributed in the QP solution. The QP algorithm results in a type of load balancing, where the load is the desired command and the algorithm balances this load among various actuators. Examination of figs. 1-2 shows that both algorithms seemed to have their sensitivities increase once one of the control surfaces became saturated. Also, the rate of change in the control surface deflections was generally smoother when using the QP algorithm, as can be seen in fig. 2.

An acceleration error for each algorithm was computed as the l_2 norm of the difference between the accelerations that result from the commanded deflections and the desired accelerations. The acceleration error is given by

$$\text{accelError}(\dot{x}_{cmd}) = \left\| Bu_{cmd} - \dot{x}_{cmd} \right\|_2 \tag{8}$$

where $\dot{x}_{cmd}$ is the vector of desired accelerations and u_{cmd} is the vector of commanded deflections determined for $\dot{x}_{cmd}$ by the algorithm. The term Bu_{cmd} gives the accelerations resulting from the commanded deflections. Figures 3-4 show that the acceleration errors were larger for the QP method for this simulation. The LP errors were essentially zero and were within the numerical noise floor when the desired accelerations are achievable. This observation means that the LP solution is the same as a two-step optimization procedure where error is minimized first and control is minimized within this solution as a secondary objective. The QP method does not result in the same zero errors. Interestingly, a difference between the LP and QP solutions is that, for ε small, the LP solution does not constitute a trade-off between performance (error) and cost (control), whereas the QP solution does. This feature may be part of the reason for the increased sensitivity of the LP algorithm. Nevertheless the QP error was small and could be made smaller by changing ε.

Figures 5-8 compare the results of computations for the ICE with a pure roll acceleration command given as the desired acceleration, *e.g.,* the pitch and yaw acceleration components were identically zero. The simulation input had the roll acceleration start at zero and step up to 1.1 times the maximum attainable roll acceleration of 351.94 deg/s^2 for the ICE model. There were 110 points in the interval. To test the sensitivity at each point in this simulation, Δ always had the roll acceleration equal to one and the pitch and yaw accelerations equal to zero. The results of the computation, as shown in figs. 5-8, tell a similar story as figs. 1-4.

An interesting observation in the plots is that both methods show increased sensitivity when the commanded accelerations approach the achievable acceleration limits. Figures 1 and 5 provide good examples of this "terminal" sensitivity for the LP and the QP methods, as the commanded pitch and roll

approach the maximum attainable limits. The terminal sensitivity is a problematic feature for both methods.

The sensitivity of the algorithms was analyzed for points randomly sampled from one of two sets: a cube containing the AMS or the boundary of the AMS. Computations using input from these two sets allow us to examine the overall sensitivity of the algorithms and the terminal sensitivity of the algorithms at the boundary of the AMS as a function of the parameter ε. A total sensitivity metric for a set of points was computed to assess the sensitivity of the algorithms to the data. The total sensitivity of a set, S, is given by

$$\text{sens}_{tot}(S) = \frac{1}{\text{size}(S)} \sum_{\dot{x}_{cmd} \in S} \frac{\left\| u_{cmd} - u_{delta} \right\|_2}{\left\| \Delta \right\|_2} \tag{9}$$

where u_{cmd} is the vector of deflections determined by the algorithm for $\dot{x}_{cmd}$, u_{delta} is the vector of deflections for $\dot{x}_{cmd} + \Delta$, with Δ a small random vector of fixed length, and size(S) is the number of elements in the set S.

Figure 9 shows the total sensitivity for ten sets of 1000 desired accelerations randomly sampled from a cube with its boundary outside the attainable moment set, see table 3 for the cube position limits. The length of the vector Δ that was used for the sensitivity calculation was 2. The total sensitivity for each of the ten sets of 1000 acceleration vectors was determined for various values of ε. Each marker in fig. 9 represents the total sensitivity for a set of 1000 accelerations using the LP or QP algorithm with the specified value of ε. From fig. 9, it can be seen that the QP algorithm was less sensitive than the LP algorithm for points randomly sampled from a cube around the AMS. Additionally, fig. 9 shows that the sensitivity of the LP algorithm is quite similar when $\varepsilon \le 10^{-4}$, whereas, the sensitivity decreases when $\varepsilon \ge 10^{-3}$. The sensitivity of the QP algorithm is relatively similar for the values of ε used in the computations, as seen in fig. 9.

A total acceleration error was computed to assess the acceleration errors for the algorithms over a set of points. The total acceleration error for a set, S, is given by

$$\text{accelError}_{tot}(S) = \frac{1}{\text{size}(S)} \sum_{\dot{x}_{cmd} \in S} \left\| B u_{cmd} - \dot{x}_{cmd} \right\|_2 \tag{10}$$

where $\dot{x}_{cmd}$ is the vector of desired accelerations, u_{cmd} is the vector of deflections determined for $\dot{x}_{cmd}$, and size(S) is the number of elements in the set S. The term Bu_{cmd} gives the accelerations resulting from the commanded deflections. When the desired accelerations were randomly sampled from the cube containing the AMS, the total acceleration error was 38.6 deg/s^2 for the LP algorithm and 32.7 deg/s^2 for the QP algorithm for each set of 1000 points. The values of the total errors were quite large because some of the points in each of the random sets were unattainable, resulting in large errors that dominated the smaller errors.

Figure 10 shows the total sensitivity for different values of ε for points randomly sampled from the boundary of the AMS. Computations showed that for accelerations on the boundary of the AMS, the QP algorithm was less sensitive than the LP algorithm when $\varepsilon \le 10^{-4}$. When $\varepsilon = 10^{-3}$, the LP and QP have comparable sensitivities. When ε was equal to 10^{-2}, the LP algorithm was generally less sensitive than the QP algorithm. The sensitivity metrics for both the LP and the QP algorithms was smaller for accelerations from the boundary of the AMS when $\varepsilon = 10^{-3}$ than when $\varepsilon \le 10^{-4}$. The acceleration errors for this simulation are given in Table 4. The total acceleration errors for both algorithms increased as ε increased, due to the increased priority given to the control minimization. These results confirm the

existence of a trade-off between performance (error) and cost (control) of both algorithms on the boundary of the AMS. Additionally, there is a trade-off between the sensitivity of the algorithms and the error when $\varepsilon \geq 10^{-4}$.

IV. Conclusion

Comparisons were made between control allocation schemes using an LP algorithm and a QP algorithm for minimizing an l_1 and an l_2 cost function, respectively. It was observed that the allocation of the control effectors was more distributed in the QP control allocation computation. Both the LP and the QP control allocation algorithms show "terminal sensitivity" or increased sensitivity when the commanded accelerations approached the achievable acceleration limits. The QP algorithm was less sensitive than the LP algorithm for points randomly sampled from a cube around the AMS. On the boundary of the AMS, both algorithms had increased sensitivity. The results confirm the existence of a trade-off between performance (error) and cost (control) for both algorithms on the boundary of the AMS and for the QP algorithm inside the AMS. Additionally, for accelerations near the boundary of the AMS, both algorithms exhibit a trade-off between sensitivity and error when $\varepsilon \geq 10^{-4}$. The results presented here suggest that designers, especially those working on next generation aircraft, would be wise to examine the sensitivity of the chosen control allocation algorithm for their application.

References

[1] "Fundamental Aeronautics Program Subsonic Fixed Wing Project Reference Document," National Aeronautics and Space Administration, May 2006, URL: http://www.aeronautics.nasa.gov/programs_fap.htm

[2] Acosta, D. M., KrishnaKumar, K., and Frost, S. A.. "Towards Intelligent Control for Next Generation Aircraft," *Proceedings International Powered Lift Conference*, 2008 Royal Aeronautical Society, London, England, July 2008.

[3] Cameron, D., Princen, N., "Control Allocation Challenges and Requirements for the Blended Wing Body," *AIAA Guidance, Navigation and Control Conference and Exhibit*, AIAA, Denver, Colorado, August 2000, AIAA-2000-4539.

[4] Page, A. B., Steinberg, M. L., "High-Fidelity Simulation Testing of Control Allocation Methods," *AIAA Guidance, Navigation and Control Conference and Exhibit*, AIAA, Monterey, California, August 2002, AIAA-2002-4547.

[5] Bodson, M., "Evaluation of Optimization Methods for Control Allocation," *AIAA Journal of Guidance, Control, and Dynamics*, vol. 25, no. 4., pp.703-711, 2002.

[6] Davidson J. B., Lallman, F. J., and Bundick, W. T., "Real-Time Adaptive Control Allocation Applied to a High Performance Aircraft," 5th SIAM Conference on Control & Its Applications, 2001.

[7] Oppenheimer, M. W. and Doman, D. B., "A Method for Including Control Effector Interactions in the Control Allocation Problem," report AFRL-VA-WP-TP-2007-309, Air Force Research Laboratory, Wright-Patterson AFB, OH 45433-7542, 2007.

[8] Nocedal, J. and Wright, S. J., *Numerical Optimization*, Springer, New York, 1999.

[9] Buffington, J., "Modular Control Law Design for the Innovative Control Effectors (ICE) Tailless Fighter Aircraft Configuration 101-3," report AFRL-VA-WP-TR-1999-3057, Air Force Research Laboratory, Wright-Patterson AFB OH 45433-7542, 1999.

[10] Luenberger, D., *Introduction to Linear and Nonlinear Programming*, Addison-Wesley, Reading, MA 1984, pp. 80.

[11] Härkegård, O., "Efficient Active Set Algorithms for Solving Constrained Least Squares Problems in Aircraft Control Allocation," *41st IEEE Conference on Decision and Control*, Vol. 2, pp. 1295-1300.

[12] Peterson, J. A. M. and Bodson, M., "Constrained Quadratic Programming Techniques for Control Allocation," *IEEE Transactions on Control Systems Technology*," Vol. 14, No. 1, January 2006, pp. 91-98.

[13] Enns, D., "Control Allocation Approaches," *AIAA Guidance, Navigation and Control Conference and Exhibit*, AIAA, Boston, MA, August 1998, AIAA-1998-4109.

[14] "Quadratic Programming Control Allocation Toolbox (QCAT)," Matlab® Central, August 2008, URL: http://www.mathworks.com/matlabcentral/fileexchange/loadFile.do?objectId=4609&objectType=file

[15] J. Buffington, "Tailless Aircraft Control Allocation, *Proc. of the AIAA Guidance, Navigation, and Control Conference*, New Orleans, LA, 1997, pp. 737-747.

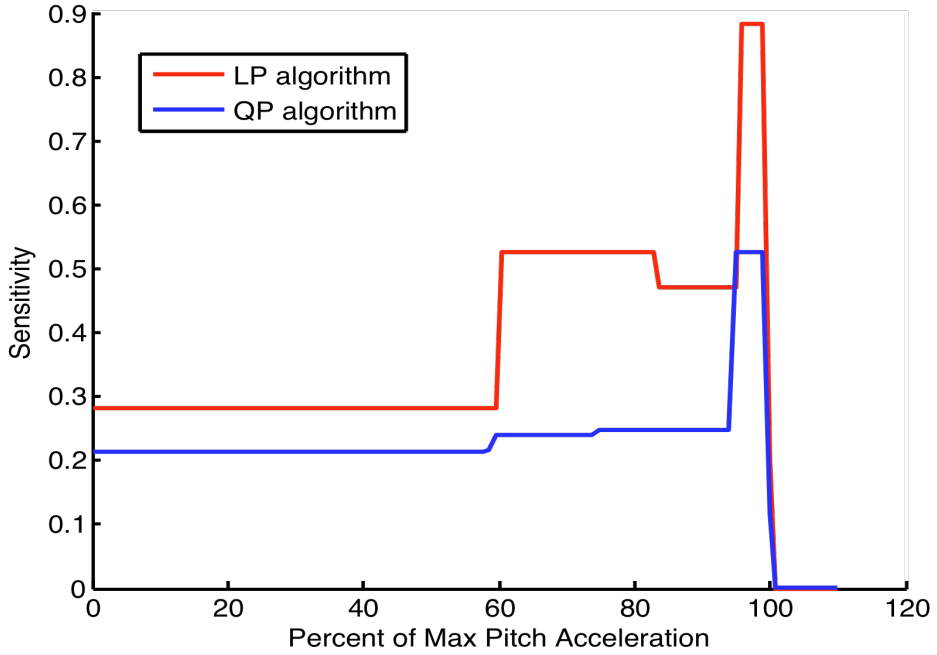

Figure 1. Sensitivity of LP versus QP algorithm for pure commanded pitch.

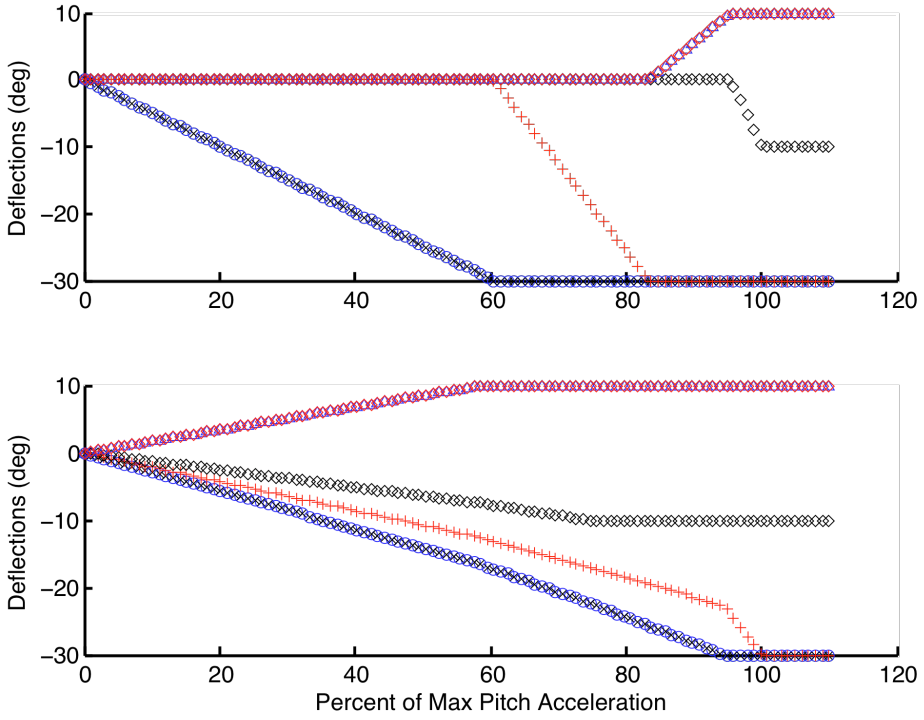

Figure 2. a) Commanded control surface deflections for pure commanded pitch from LP algorithm and b) QP algorithm.

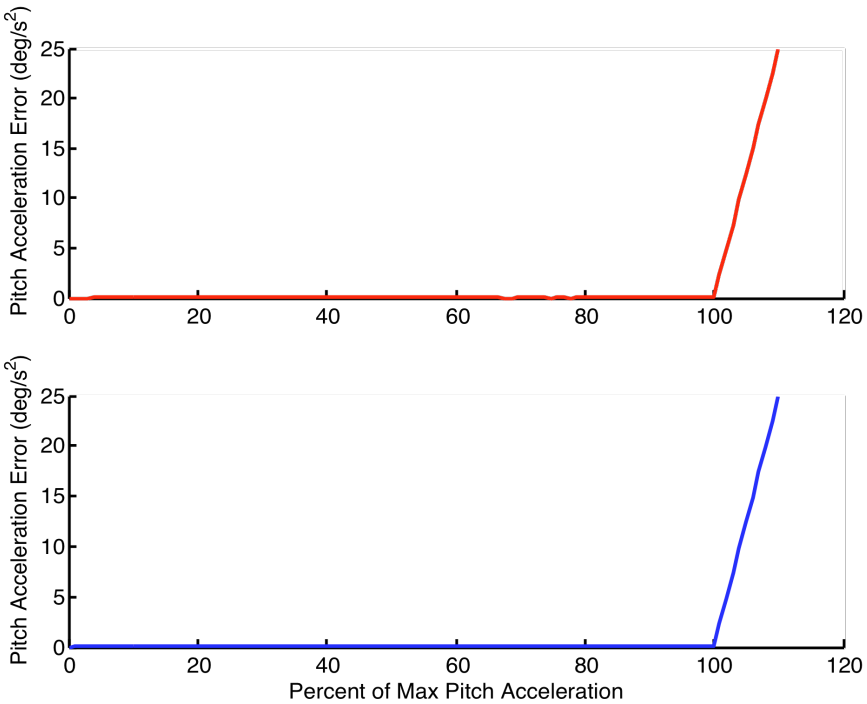

Figure 3. a) Acceleration errors for pure commanded pitch from LP algorithm and b) QP algorithm.

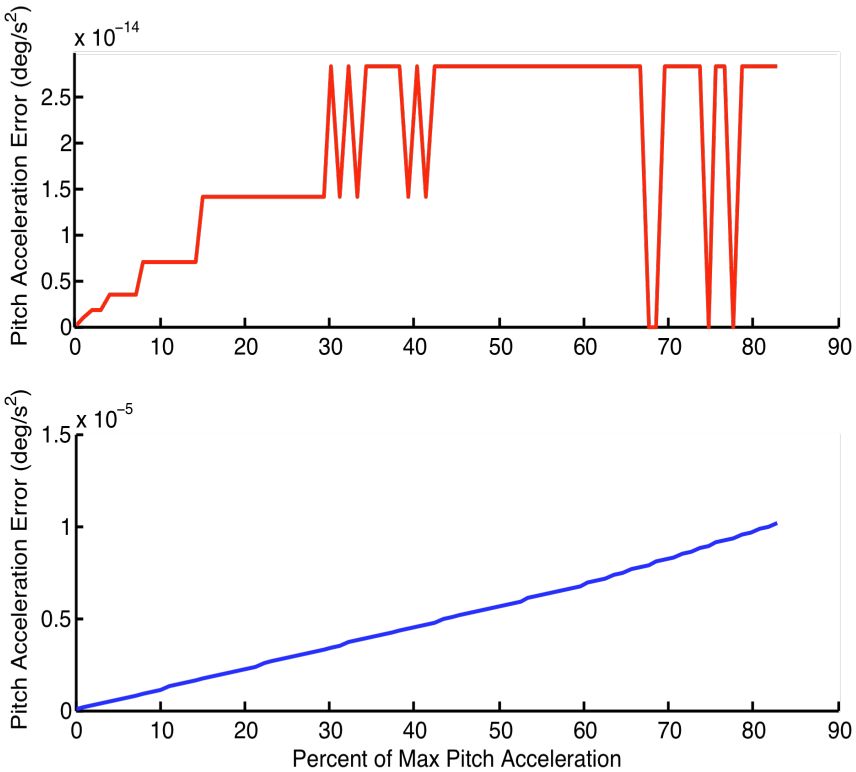

Figure 4. a) Detail of acceleration errors for pure commanded pitch from LP algorithm and b) QP algorithm.

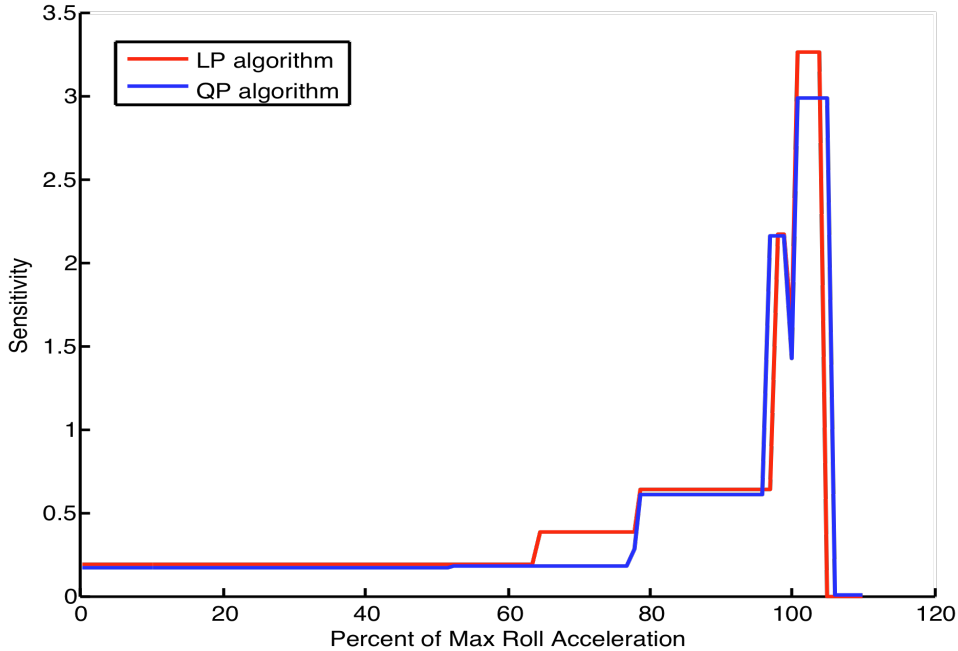

Figure 5. Sensitivity of l1 versus l2 algorithm for pure commanded roll.

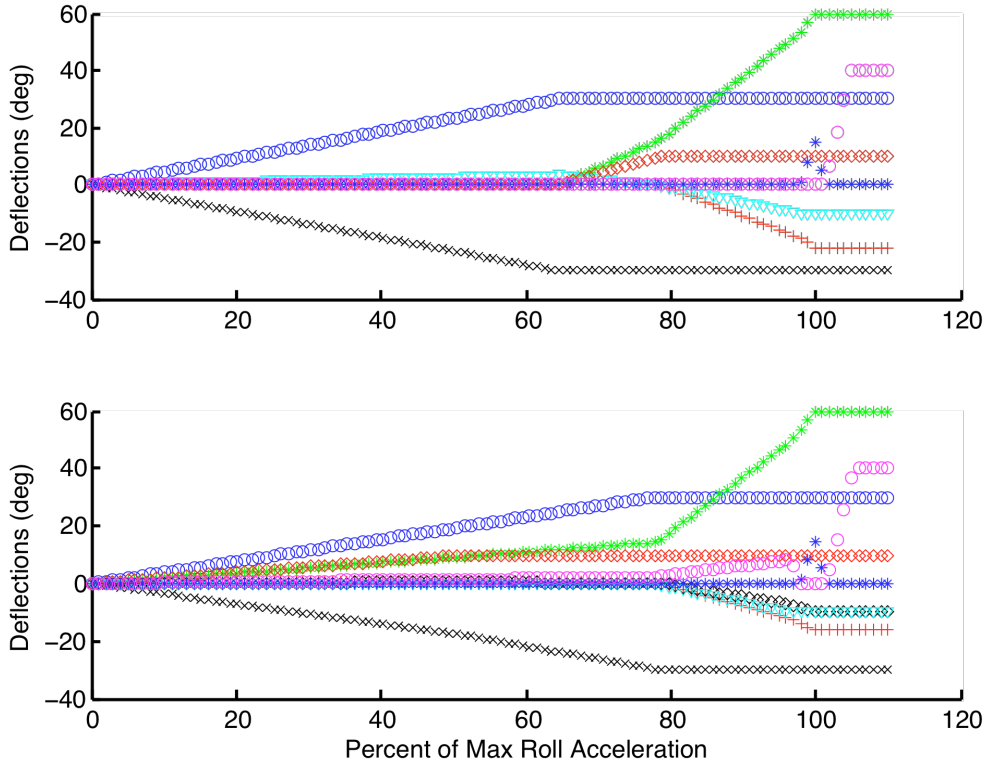

Figure 6. a) Commanded control surface deflections for pure commanded roll from LP algorithm and b) QP algorithm.

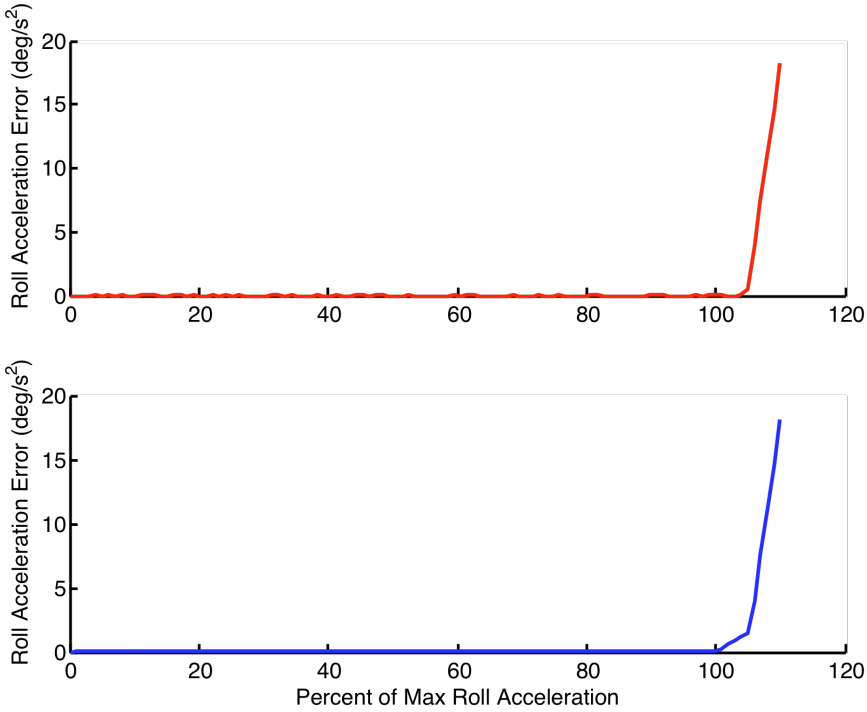

Figure 7. a) Acceleration errors for pure commanded roll from LP algorithm and b) QP algorithm.

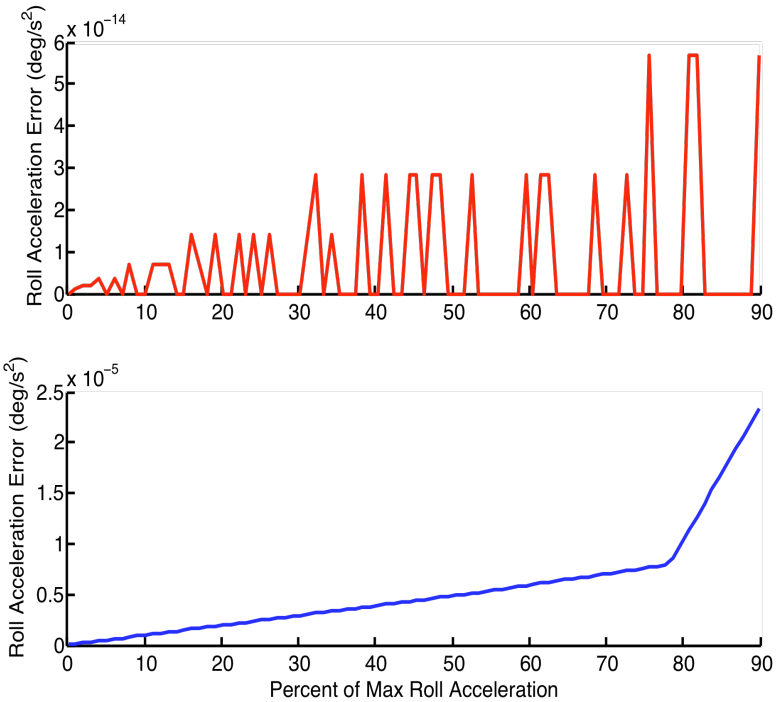

Figure 8. a) Detail of acceleration errors for pure commanded roll from LP algorithm and b) QP algorithm.

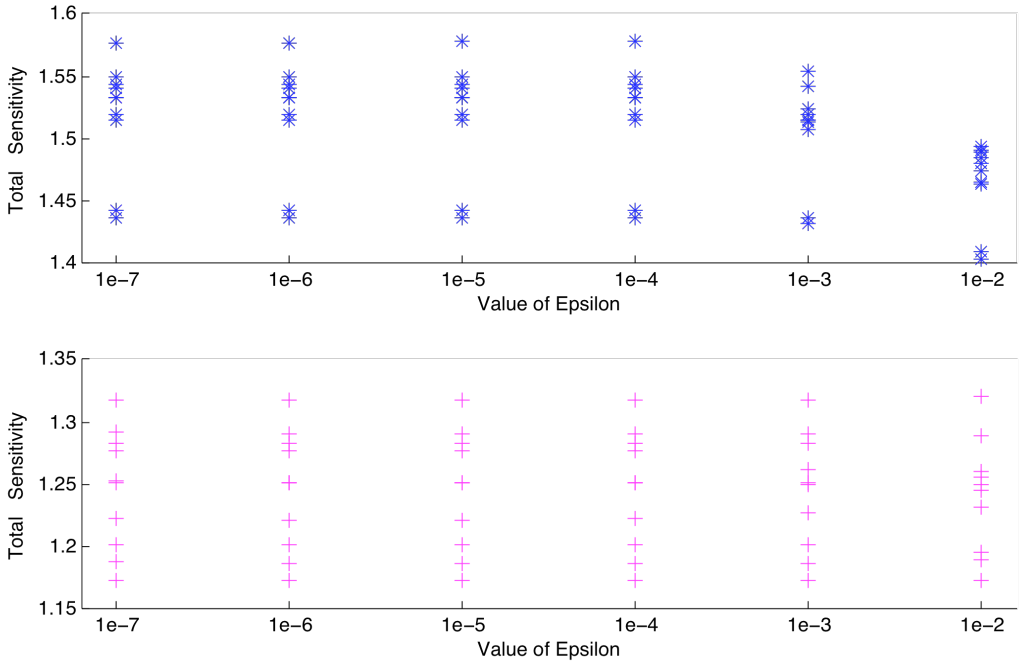

Figure 9. a) Total sensitivity for sets of 1000 random desired accelerations in cube around AMS for LP algorithm and b) QP algorithm.

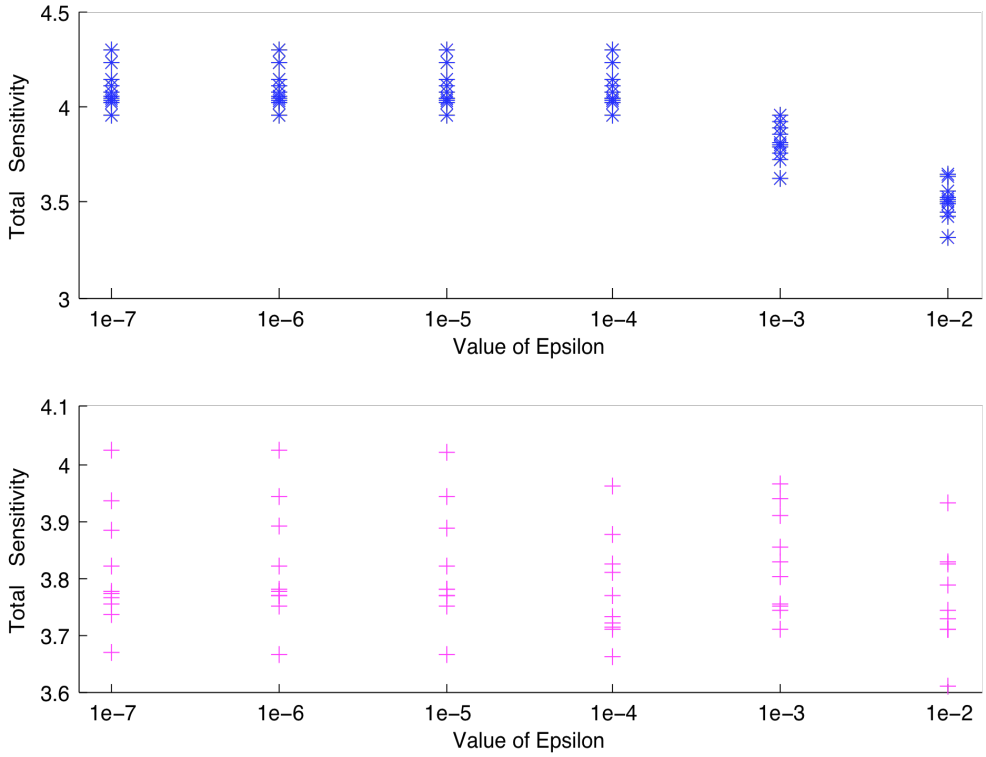

Figure 10. a) Total sensitivity for sets of 1000 random desired accelerations on boundary of AMS for LP algorithm and b) QP algorithm.

	left elevon	right elevon	pitch flaps	left all-moving tip	right all-moving tip	pitch thrust vector-ing	yaw thrust vector-ing	left spoiler slots	right spoiler slots	left leading-edge flaps	right leading-edge flaps
pitch rate	-2.5114	-2.5115	-1.9042	-0.9494	-0.9494	-1.1329	0	1.5046	1.5046	-0.0003	-0.0004
roll rate	3.7830	-3.7830	0	1.8255	-1.8255	0	-0.0790	-2.0956	2.0957	-0.3067	0.3067
yaw rate	0.0453	-0.0453	0	-0.2081	0.2081	0	-0.8038	-0.0283	0.0283	0.0937	-0.0937

Table 1. *B* matrix entries.

	left elevon	right elevon	pitch flaps	left all-moving tip	right all-moving tip	pitch thrust vector-ing	yaw thrust vector-ing	left spoiler slots	right spoiler slots	left leading-edge flaps	right leading-edge flaps
u_{min} (deg)	-30	-30	-30	0	0	-10	-10	0	0	0	0
u_{max} (deg)	30	30	30	60	60	10	10	10	10	40	40

Table 2. Minimum and maximum control surface deflection limits (deg).

	Pitch	Roll	Yaw
Acceleration Limits (deg/s^2)	±333.100	±351.941	±25.772

Table 3. Limits (deg/s^2) for cube around AMS.

	Value of ε					
Algorithm	1e-2	1e-3	1e-4	1e-5	1e-6	1e-7
LP	9.64e-3	3.472e-4	1.794e-8	1.794e-8	6.362e-9	3.80e-13
QP	5.162e-2	1.607e-3	9.082e-5	1.177e-6	2.560e-8	8.048e-10

Table 4. Total acceleration errors (deg/s^2) for simulations where input sets contain random desired accelerations from the boundary of the AMS.