

Peak-Seeking Optimization of Trim for Reduced Fuel Consumption: Architecture and Performance Predictions

Jacob Schaefer¹ and Nelson A. Brown²

NASA Dryden Flight Research Center, Edwards, California, 93523

A peak-seeking control approach for real-time trim configuration optimization for reduced fuel consumption has been developed by researchers at the National Aeronautics and Space Administration (NASA) Dryden Flight Research Center to address the goals of the NASA Environmentally Responsible Aviation project to reduce fuel burn and emissions. The peak-seeking control approach is based on a steepest-descent algorithm using a time-varying Kalman filter to estimate the gradient of a performance function of fuel flow versus control surface positions. In real-time operation, deflections of symmetric ailerons, trailing-edge flaps, and leading-edge flaps of an F/A-18 airplane (McDonnell Douglas, now The Boeing Company, Chicago, Illinois) are controlled for optimization of fuel flow. This paper presents the design and integration of this peak-seeking controller on a modified NASA F/A-18 airplane with research flight control computers. A research flight was performed to collect data to build a realistic model of the performance function and characterize measurement noise. This model was then implemented into a nonlinear six-degree-of-freedom F/A-18 simulation along with the peak-seeking control algorithm. With the goal of eventual flight tests, the algorithm was first evaluated in the improved simulation environment. Results from the simulation predict good convergence on minimum fuel flow with a 2.5-percent reduction in fuel flow relative to the baseline trim of the aircraft.

Nomenclature

Acronyms

KCAS	=	knots calibrated airspeed
KF	=	Kalman filter
NASA	=	National Aeronautics and Space Administration
NDI	=	nonlinear dynamic inversion
PE	=	persistent excitation

Symbols

b	=	gradient estimate
F	=	performance measurement matrix
f	=	performance parameter
G	=	gain applied to gradient
H	=	observation matrix
I	=	identity matrix
K	=	Kalman filter gain
k	=	iteration sample
M	=	number of previous samples to use in Kalman filter
n	=	number of control surfaces used for optimization
O	=	higher-order terms
P	=	state covariance matrix
Q	=	process noise covariance matrix
R	=	measurement noise covariance matrix
u	=	control surface commands
v	=	measurement noise

¹ Aerospace Engineer, Flight Controls and Dynamics Branch, P.O. Box 273 / Mail Stop 4840D, jacob.schaefer@nasa.gov, AIAA Member.

² Aerospace Engineer, Flight Controls and Dynamics Branch, P.O. Box 273 / Mail Stop 4840D, nelson.brown@nasa.gov, AIAA Member.

w	=	process noise
x	=	surface deflections
z^{-1}	=	unit delay
ζ	=	Kalman filter states
$\bar{\cdot}$	=	vector notation
$\hat{\cdot}$	=	predicted
\cdot^T	=	transpose

I. Introduction

The NASA Environmentally Responsible Aviation project is seeking to mitigate the impact of aviation on the environment by reducing fuel consumption and emissions from aircraft.¹ In 2011, United States air carriers, for domestic flights, consumed 12.1 billion gallons of fuel, releasing 114.6 million metric tons of carbon dioxide (CO₂) equivalent into the atmosphere.^{2,3} Small-percentage savings in fuel consumption rates for commercial aircraft will yield both large cost savings for airlines and reductions in emissions. This paper presents a control algorithm with the potential to reduce fuel consumption during cruise flight. The system performs real-time online optimization of the trim configuration of the secondary longitudinal control effectors of the aircraft, such as flaps or symmetric aileron deflections, to minimize fuel flow to the engines while the stabilators maintain pitch trim. The algorithm is based on the techniques of Ryan and Speyer (Ref. 4) adapted for the trim optimization problem.

While commercial aircraft could benefit greatly from this technology, it is not aircraft type-specific and could be implemented on a variety of other aircraft, ranging from fighter jets to high-aspect-ratio unmanned aircraft. With the goal of raising the technology readiness level of this control algorithm and eventually performing flight-testing, a simulation study was performed to investigate implementing the algorithm on the X-48B Blended Wing Body research aircraft (The Boeing Company, Chicago, Illinois). Due to the limitations listed in Ref. 5, another platform was chosen for flight research and testing: a NASA F/A-18 airplane, tail number 853, with research flight control computers. In this paper, the design and integration of the peak-seeking algorithm onto the NASA F/A-18 airplane is presented along with simulation results. Early in the development stage, it was found desirable for the purpose of analysis to have more accurate simulation models for the drag and pitching moment created by the secondary pitch surfaces, as well as models for the fuel flow sensor noise characteristics. Thus, an early research flight was flown to collect data to improve the simulation models. The control algorithm was then evaluated in the improved simulation environment. The results from the simulation tests and the expected reduction in fuel flow are presented. The peak-seeking control algorithm was then flown on five research flights; the results are given in Ref. 6.

II. Background

In cruise flight, aircraft typically use a single effector, often through the horizontal stabilator or elevator, to maintain pitch trim control. Other surfaces capable of producing pitching moments, such as symmetric ailerons, trailing-edge flaps, and leading-edge flaps, can be used in combinations to trim the aircraft for steady level flight. Further constraints can then be placed in order to choose the combination that maximizes performance. One of the first programs to investigate the use of multiple control surfaces for enhanced performance was the Mission Adaptive Wing (MAW) program, using an F-111 (General Dynamics Corporation, Falls Church, Virginia). Through the use of the leading-edge flaps and multiple trailing-edge control surfaces, the MAW program demonstrated the ability to achieve improvements in lift-to-drag ratio (L/D) in the range of approximately 10 to 20 percent at typical cruise conditions.^{7,8} For transport class aircraft, a study by Airbus (Toulouse, France) in the 1980s showed a 3- to 9-percent increase in L/D during cruise by varying the wing shape through the use of variable camber control.⁹ Based on a similar concept, in 2009 The Boeing Company issued a performance improvement package for their 777 airplane that included a software upgrade to change the trim aileron deflection in cruise flight from 0 deg to 2 deg trailing-edge down to achieve a spanwise lift distribution that is more efficient.¹⁰ While this is a fixed deflection over all flight conditions, some modern transport aircraft, such as the Boeing 787, use extensive lookup tables for control surface positions based on flight condition, flight phase, and aircraft loading in order to determine the control surface positions for the best performance.¹¹ Models for lookup tables, however, can require extensive testing and research to cover the entire flight envelope and all possible aircraft configurations, yet still may contain large amounts of uncertainty. Variations between aircraft as well as variations that develop over the lifetime of an airframe may add error to the models. Another class of aircraft in which aerodynamic efficiency is critical and many control surfaces can often be used for pitch trim is that of high-altitude, long-endurance aircraft. These aircraft typically have high fuel fractions and thus often see large changes in lift coefficient over cruise flight, requiring significant re-trimming. Reference 12 presents a study that predicted that a 5- to 15-percent increase in range for

such aircraft is achievable by scheduling multiple trailing-edge control surface trim deflections for minimum drag across the entire cruise range. This method used analytic and computational methods to determine the optimal control surface positions a priori. With the goal of real-time optimization, without the need for a priori information, research by Gilyard et al. at the NASA Dryden Flight Research Center used techniques to find the minimum drag by modulating symmetric aileron deflections periodically and estimating the resulting drag in real time on an L-1011 airplane (Lockheed Corporation, now Lockheed Martin, Bethesda, Maryland).¹³⁻¹⁵

This paper presents a real-time optimization algorithm for determining the optimal configuration without the need for any performance models and that can optimize any number of control effectors simultaneously. It can also utilize performance measurements from onboard sensors, rather than from performance models or estimations of drag.

III. Architecture

The peak-seeking control algorithm presented here has a time-varying Kalman filter (KF) that utilizes in-flight measurements as a basis for the optimization. This algorithm has been analyzed in other previous peak-seeking applications.^{4, 16} Here it is applied to trim optimization. Using this algorithm, a trim solution can be found for a set of control surfaces to minimize a performance function defined as fuel flow versus trim configuration.

A. Peak-Seeking Algorithm Design

The peak-seeking control method is shown in Fig. 1 and was adapted from Ryan and Speyer.^{4,19} The peak-seeking algorithm assumes that stabilization is achieved through inner-loop control. Outer-loop trim control optimizes the measurable performance parameter, f , which in this application is measured as the rate of fuel flow to the engines. The performance parameter is assumed, under constant flight condition, to be solely a function of \bar{x} which is a vector of measurements of n independent control surface positions which are to be optimized. The function $f(\bar{x}_k)$ is defined as the performance function which is unknown globally but is measurable. At each iteration of the optimization, differences are calculated between the previous and current surface positions, $\bar{x}_{k-1} - \bar{x}_k$, and the previous and current performance parameter measurements, $f_{k-1} - f_k$. From these differences, a linear time-varying Kalman filter provides estimates of the current gradient vector, \bar{b}_k , the change in performance per change in surface positions. In this application, a constant negative gain multiplier, G , is applied to the gradient and used to drive the plant in a steepest descent toward the minimum fuel flow as given in Eq. (1) where the commanded trim surface positions is given as \bar{r}_k .

$$\bar{r}_k = (G \bar{b}_k) + \bar{r}_{k-1} \quad (1)$$

A persistent excitation signal (PE_k) is also added to the estimated control command, \bar{r}_k , to provide observability of the performance function to the Kalman filter.

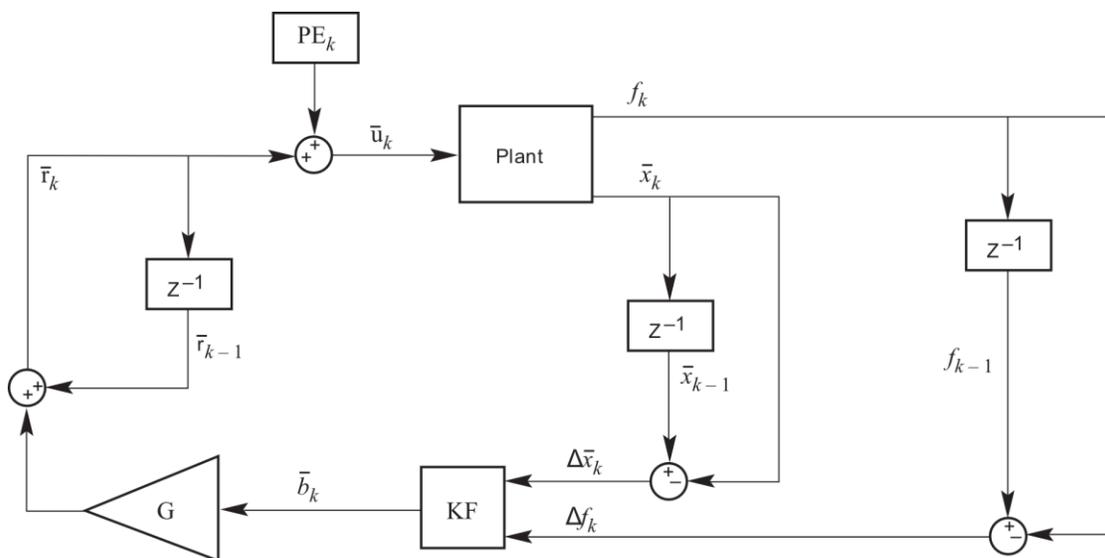


Figure 1. The peak-seeking control architecture.

B. Kalman Filter Design

The heart of this peak-seeking control scheme is the use of the linear time-varying Kalman filter, which uses measurements of the performance function and control surface positions, and whose states contain the current gradient estimate which is formulated as follows:

A Taylor series expansion of the performance function $f(\bar{x})$ about \bar{x}_k is given by Eq. (2):

$$f(\bar{x}) \approx f(\bar{x}_k) + \bar{b}_k^T (\bar{x} - \bar{x}_k) + O(\bar{x} - \bar{x}_k) \quad (2)$$

where $O(\cdot)$ represents higher-order terms. Assuming the performance function can be sufficiently treated as linear at any combination of control surface positions, the difference between the previous and current performance function magnitude can then be approximated as shown in Eq. (3):

$$\Delta f_k = \bar{b}_k^T \Delta \bar{x}_k \quad (3)$$

where $\Delta f_k = f(\bar{x}_{k-1}) - f(\bar{x}_k)$ and $\Delta \bar{x}_k = \bar{x}_{k-1} - \bar{x}_k$. This is further denoted for n effectors in Eq. (4).

$$\Delta \bar{x}_k = \begin{bmatrix} \Delta x_{1k} \\ \Delta x_{2k} \\ \vdots \\ \Delta x_{nk} \end{bmatrix} = \begin{bmatrix} x_{1k-1} - x_{1k} \\ x_{2k-1} - x_{2k} \\ \vdots \\ x_{nk-1} - x_{nk} \end{bmatrix} \quad (4)$$

Equation (3) can be expanded and rewritten using Eq. (4) to give Eq. (5):

$$\Delta f_k = \begin{bmatrix} b_{1k} \\ b_{2k} \\ \vdots \\ b_{nk} \end{bmatrix}^T \begin{bmatrix} \Delta x_{1k} \\ \Delta x_{2k} \\ \vdots \\ \Delta x_{nk} \end{bmatrix} \quad (5)$$

where f , x_1 , x_2 , ..., x_n are measureable quantities and the gradient \bar{b} is the only unknown and is to be estimated. Since the gradient may change with \bar{x} , and measurements of \bar{x}_k , and f_k may be noisy, a Kalman filter is an appropriate choice for an estimator. The Kalman filter states are chosen to be the gradient vector, given in Eq. (6).

$$\zeta_k = \begin{bmatrix} b_{1k} \\ b_{2k} \\ \vdots \\ b_{nk} \end{bmatrix} \quad (6)$$

Equations (2) through (5) are defined by using only the current and previous measurements, k and $k - 1$. For better estimation of the gradient, any number of prior measurements, M , can be utilized by the Kalman filter. The measurement variables in Eq. (5) can then be expanded by defining two new matrices, ΔF_k , and H_k as follows in Eqs. (7) and (8):

$$\Delta F_k = \begin{bmatrix} f(\bar{x}_{k-1}) - f(\bar{x}_k) \\ f(\bar{x}_{k-2}) - f(\bar{x}_k) \\ \vdots \\ f(\bar{x}_{k-M}) - f(\bar{x}_k) \end{bmatrix}^T \quad (7)$$

$$H_k = \begin{bmatrix} x_{1k-1} - x_{1k} & x_{2k-1} - x_{2k} & \dots & x_{nk-1} - x_{nk} \\ x_{1k-2} - x_{1k} & x_{2k-2} - x_{2k} & \dots & x_{nk-2} - x_{nk} \\ \vdots & \vdots & \ddots & \vdots \\ x_{1k-M} - x_{1k} & x_{2k-M} - x_{2k} & \dots & x_{nk-M} - x_{nk} \end{bmatrix} \quad (8)$$

The measurement equation of the linear time-varying Kalman filter then takes the form of Eq. (9):

$$\Delta F_k = \zeta_k^T H_k^T + v_k \quad (9)$$

where ΔF_k and H_k are composed of measurable quantities and v_k represents a zero-mean Gaussian white-noise process with measurement covariance matrix R_k . Given the unknown true dependence of the performance function on surface positions \bar{x}_k , the state is modeled as a Brownian noise process and the linear time-varying Kalman filter process equation is then given by Eq. (10):

$$\zeta_k = \zeta_{k-1} + w_k \quad (10)$$

where w_k represents a zero-mean Gaussian white-noise process with covariance matrix Q_k . The linear time-varying Kalman filter implementation is then given from Eqs. (11) through (15):

$$K = \hat{P}_k H_k^T (H_k \hat{P}_k H_k^T + R_k)^{-1} \quad (11)$$

$$\zeta_k = \hat{\zeta}_k + K(\Delta F_k - H_k \hat{\zeta}_k) \quad (12)$$

$$P_k = (I - KH_k) \hat{P}_k \quad (13)$$

$$\hat{\zeta}_{k+1} = \zeta_k \quad (14)$$

$$\hat{P}_{k+1} = P_k + Q_k \quad (15)$$

where P_k and $\hat{P}_{(\cdot)}$ are current and predicted state covariance matrices, respectively, at the corresponding iterations. Similarly, ζ_k and $\hat{\zeta}_{(\cdot)}$ are the current and predicted Kalman filter state estimates which correspond to the gradient of the performance function, \bar{b}_k . The noise covariance matrices R_k and Q_k are used as tuning parameters influencing the performance of the Kalman filter, detailed below.

C. Persistent Excitation

The purpose of the persistent excitation signal is for observability of the performance function. For some systems, random disturbances may be sufficient; however, a methodical approach was implemented here.

A method was chosen for good observability across all axes and involved adding an excitation signal that was helical about the commanded trajectory. This was done by rotating a vector by $\frac{n}{360}$ deg, with respect to the previous excitation command, around a circle in the orthogonal plane to the trajectory. The radius of the circle was scaled as a linear function of the magnitude of the trajectory step, $\bar{r}_{k-1} - \bar{r}_k$. Conceptually this is shown for a two-effector case and a three-effector case in Fig. 2. For the two-effector case, the addition of the persistent excitation signal results in a sinusoidal trajectory; the three-effector case results in a spiral. This method can be expanded upon to include additional dimensions for any number of effectors, n .

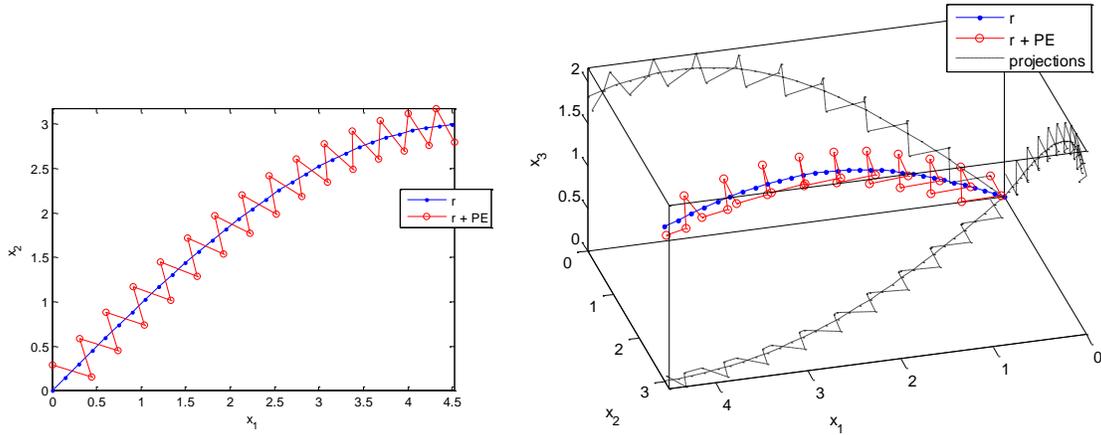


Figure 2. Surface commands over 25 iterations with and without persistence excitation for two effectors (left), and three effectors (right).

D. Initial Excitation

The Kalman filter must initially be seeded with a set of surface deflections and performance measurements. This derives from examining Eqs. (7) and (8), where ΔF_k and H_k are computed with a series of previous sampled measurements, M , that are required to estimate the gradient. The matrices can be seeded with zeros, however, to begin with a good initial trajectory, an initial excitation routine should be performed.

The method of initially seeding these matrices with measurement data was chosen to give good observability of the performance function in each axis resulting in a good initial gradient estimation. For the two-effector case, this meant measuring the performance function at M evenly-distributed surface deflection points on a circle centered at the initial condition with a radius of 1 deg. For the three-effector case this involved distributing M points about a sphere of radius 1 deg.

IV. F/A-18 Implementation

A control law utilizing this peak-seeking control algorithm has been developed for the NASA Full-Scale Advanced Systems Testbed (FAST) airplane, tail number 853. The FAST airplane is a modified F/A-18 airplane with a quad-redundant research flight control system that provides health and envelope checking for two faster dual-redundant computers, the Airborne Research Test System, which are capable of housing advanced control laws written in C code or Simulink® (The MathWorks, Natick, Massachusetts) autocode for rapid design and flight testing.¹⁷ The control surfaces and engines are that of a standard F/A-18 airplane. The airplane, shown in Fig. 3, has the ability to move the leading-edge flaps (+leading edge down), trailing-edge flaps (+trailing edge down), and ailerons symmetrically (+trailing-edge down) to create pitching moment changes and act as longitudinal control effectors, in addition to the horizontal stabilators. The rudders can also create a pitching moment when deflected both in or out; however, for this experiment only the wing control surfaces were used as part of the optimization. Although the F/A-18 airplane has multiple fuel tanks, the research control system is incapable of transferring fuel between them; for transport class aircraft, however, this would likely be an important effector. The F/A-18 airplane also has fuel flow meters and temperature sensors in the fuel lines to each engine which yield a mass flow rate which is the source for the measurement of the performance function in the peak-seeking algorithm.

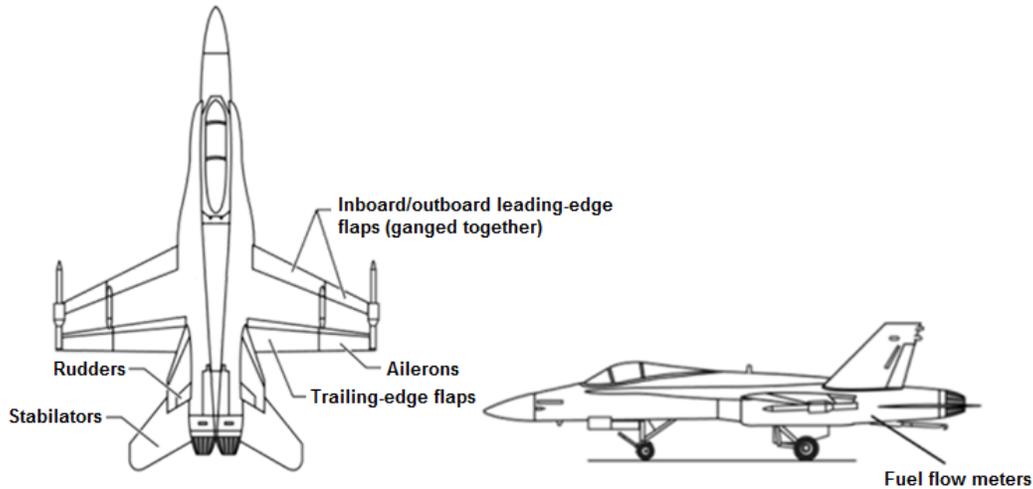


Figure 3. The F/A-18 airplane.

The peak-seeking algorithm is integrated with a NASA-designed nonlinear dynamic inversion (NDI) inner-loop control law as shown in Fig. 4.¹⁸ For steady level trimmed flight, the NDI solves for a trim state that gives zero pitch rate by use of the horizontal tail, with the leading-edge flaps and trailing-edge flaps being scheduled with angle of attack. For this application, the NDI has been modified to remove the leading-edge flap and trailing-edge flap schedules and replaced with constant zeros so that the flap positions are solely commanded from the peak-seeking controller. The outer-loop autopilot consists of a wing leveler, an altitude hold, and a velocity hold through lateral stick, longitudinal stick, and throttle position, respectively.

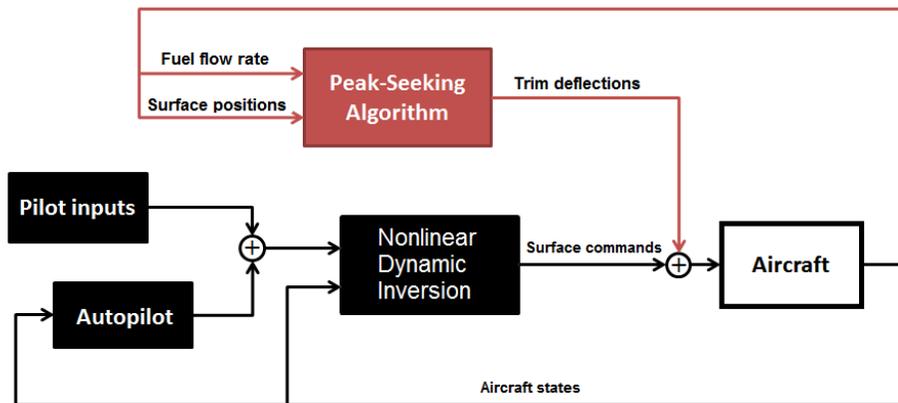


Figure 4. The peak-seeking algorithm (red) integrated with the control laws of the F/A-18 airplane.

In this configuration, the peak-seeking controller makes adjustments to the leading-edge flaps, trailing-edge flaps, and symmetric ailerons, while the inner-loop control law and the autopilot move the stabilators, throttles, differential rudders, and differential ailerons to maintain flight condition. The resulting total change in airplane drag is detected by a change in fuel flow, which is fed back to the peak-seeking algorithm.

Between iterations of the peak-seeking algorithm, the airplane must be allowed to reacquire the original flight condition and sufficient time must be allowed for the engine controller to reach a steady-state fuel flow. For small changes in surface deflections on the order of a couple degrees this was found to be in the range of 20 to 30 s. To reduce the effect of signal noise and disturbances, a low-pass filter was applied to the fuel flow prior to being used by the peak-seeking controller.

During iteration steps, a rate limit of the peak-seeking trim control commands was implemented at 2 deg/s to minimize any noticeable effects on the pilot. Also, a saturation range was added to the trim deflections that allowed for commands between 0 and 20 deg for flaps, and between +/-20 deg for ailerons, to prevent large excursions from nominal trim.

V. Simulation Models from Flight Data

The true nature of the performance function is difficult to predict analytically, and the signal noise characteristics in a true flight environment are unknown. The existing simulation models for the drag due to control surfaces are not modeled to the fidelity and accuracy required for fine optimization problems. For these reasons, a research flight was performed to collect data to build a model of the performance function. This model was then used to evaluate the feasibility of the peak-seeking algorithm in simulation prior to being flight-tested. Performance models are not needed by the algorithm but are used to generate a more realistic plant model for simulation testing. To generate the model of the performance function, a matrix of combinations of surface deflections of symmetric ailerons, trailing-edge flaps, and leading-edge flaps were commanded during a research flight and the measurement of the resulting fuel flow was recorded. A total of 80 test points were evaluated in a single flight. Each test point was performed by deflecting the three sets of control surfaces (leading-edge flaps, trailing-edge flaps, and symmetric ailerons) to a specified point and holding them there for between 30 and 60 s to measure the steady-state fuel flow. All test points were carried out with the autopilot maintaining steady level flight at a single flight condition of an altitude of 25,000 ft and a speed of 240 KCAS. Tares were performed at several points during the flight to de-trend the fuel flow signal due to reduced aircraft weight as fuel was burned off.

Second-order polynomials were fitted to the flight data across the three control effectors and were shown to yield a good fit. Figure 5 shows a three-dimensional contour plot of the approximated performance function: blue is minimum fuel flow and red is maximum fuel flow. The minimum fuel flow is at -3.1 percent from the baseline trim with symmetric ailerons at 5 deg, trailing-edge flaps at 3 deg, and leading-edge flaps at 7 deg. This condition compares to the baseline trim of the airplane, which is 0 deg symmetric aileron and flaps varying with angle of attack; however, at a medium fuel weight the trailing-edge flaps are at approximately 6 deg and the leading-edge flaps are at approximately 5 deg.

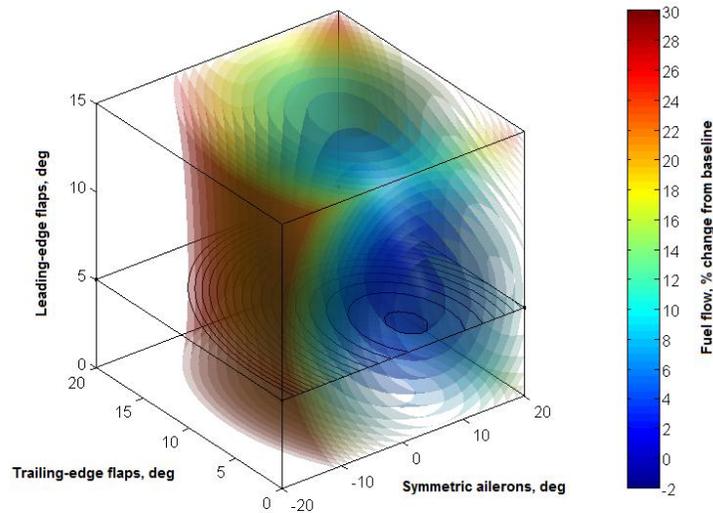


Figure 5. Contours of the performance model based on flight data.

A more detailed sampling during the research flight was performed for fuel flow changes due to just trailing-edge flaps and symmetric ailerons, with the leading-edge flaps fixed at the nominal baseline trim value of 5 deg. This was done because it was hypothesized that the trailing-edge flaps and symmetric ailerons would have a larger effect on fuel flow, essentially controlling the spanwise lift distribution. The more detailed data set collected at fixed leading-edge flaps, forms one horizontal slice in Fig. 5. Figure 6 shows a fit to this denser data set. The minimum point for this data set has a fuel flow of -2.3 percent relative to the baseline trim.

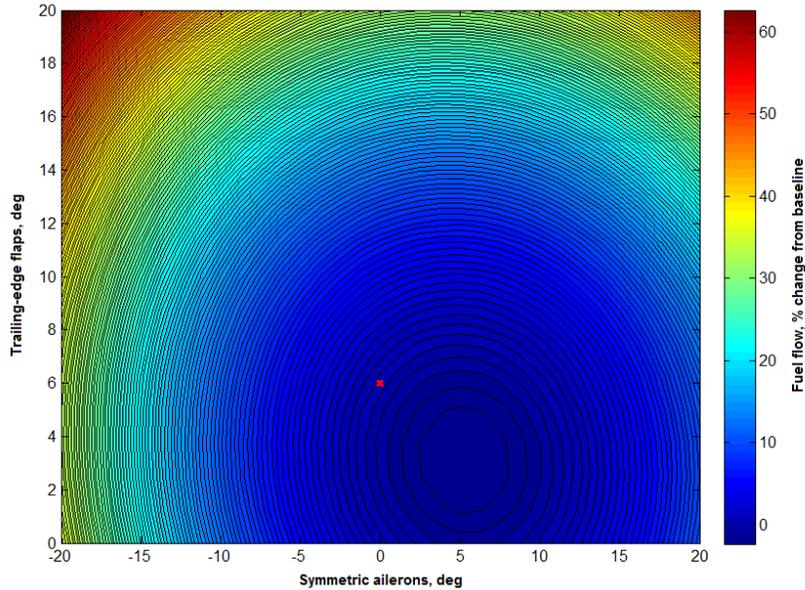


Figure 6. The performance function fit to trailing-edge flaps and aileron deflections, the baseline nominal trim for the F/A-18 airplane is denoted by a red x.

While a polynomial fits the data well, each test point had to be held for 30 to 60 s to realize the steady-state value because of the time it took for the engine controller to reach steady state and to collect enough data to filter the effects of signal noise and disturbances. The largest cause of disturbances was small changes in the atmosphere, primarily disturbances in dynamic pressure. While the pilot stated he felt no turbulence during the flight, the atmosphere is not perfectly uniform, which causes small but noticeable autothrottle response and thus affects fuel flow on the scale that is being examined here.

In order to reduce the effects of signal noise and disturbances, a tradeoff exists between using a low-pass filter on fuel flow with longer duration dwell periods between iterations of the peak-seeking controller versus using more data points at a higher frequency with shorter duration dwell and thus relying more on the Kalman filter to reject measurement noise. This tradeoff is investigated below by tuning the low-pass filter and the frequency at which the Kalman filter is iterated.

To build a model of the fuel flow signal noise and disturbances, the flight data were de-trended by the polynomial fit of the performance function calculated above to generate a zero mean disturbance signal of approximately 2000 s long. This signal was then added to the performance data to build a realistic model for testing the algorithm.

VI. Simulation

The peak-seeking controller was implemented in a six-degree-of-freedom nonlinear F/A-18 simulation along with the model of the performance function, which included the disturbance and noise signal generated from the flight data. All simulation tests were run at an altitude of 25,000 ft and a speed of 240 KCAS to match the flight condition associated with the plant model. It should be noted, however, that the peak-seeking algorithm operates independent of flight condition, assuming that signal noise and disturbance variations are small between different flight conditions or that the Kalman filter is robust to these variations.

A. Tuning in Simulation

The peak-seeking algorithm has several parameters that can be adjusted: values for the gradient gain, G ; the number of prior measurements, M ; matrices R and Q ; and the prefiltering time constant or “dwell time.” Many of the settings affect each other, and the tuning process is iterative. In this section, the effect of each tunable parameter is examined, and for simplicity only the trailing-edge flaps and ailerons are optimized with the leading-edge flap fixed at 5 deg. The same process is used for the three-surface optimization problem using ailerons, trailing-edge flaps, and leading-edge flaps. The performance of the peak-seeking algorithm is evaluated during tuning using the

following four metrics: convergence speed, minimum proximity and wander, disturbance response, and overshoots of the minimum.

It was discovered the peak-seeking controller gain, G , applied to the gradient output from the Kalman filter, had one of the largest effects on convergence to the minimum fuel flow. Figure 7 shows the resulting trajectories for a range of gain values, each run for 36 iterations. The high gain setting tended to overshoot the minimum while the low gain took a larger number of iterations and longer time to reach the minimum. The medium gain generally had good convergence on the minimum. The initial excitation circle can be seen at the start and the sawtooth-shaped pattern is the result of the persistent excitation signal.

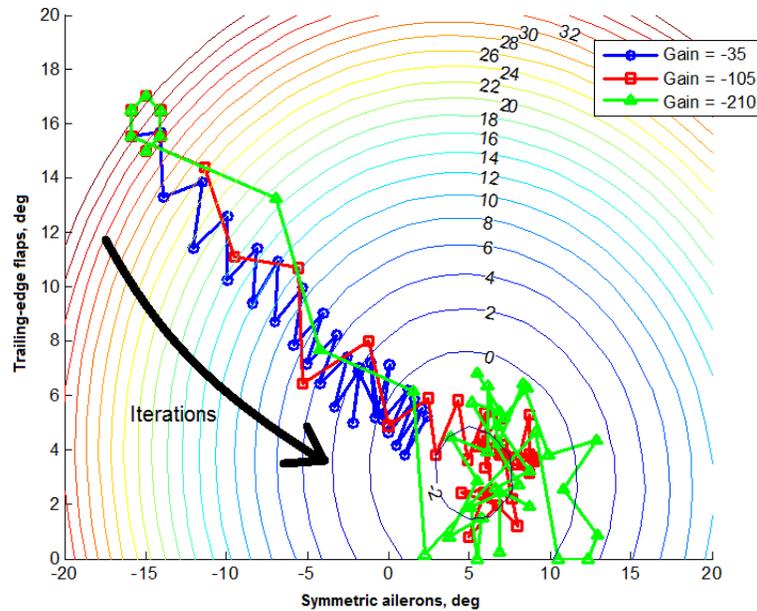


Figure 7. Peak-seeking algorithm simulation runs with varying gain value.

The number of past measurement points to be used in the Kalman filter, M , can be varied to give differing algorithm performance, as shown in Fig. 8, with all other parameters remaining the same. The time between points was 20 s, during which a moving point mean was applied to the fuel flow signal. With only two previous measurements plus the current measurement, the algorithm is clearly sensitive to disturbances and the trajectory is more erratic and deviates from the steepest-descent trajectory. With five measurements, the trajectory was generally close to the steepest descent. Using eight prior measurements the trajectory lagged the steepest descent; however, once converged on the minimum, it was less affected by noise and disturbances and remained near the minimum to a higher degree of tolerance, with less wander.

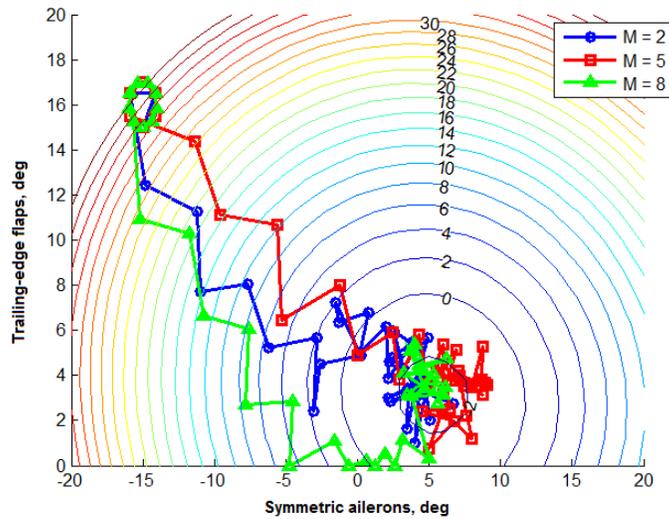


Figure 8. Varied number of measurements.

To reduce the sensitivity to disturbances for the two-measurement case and improve the eight-measurement case, the tradeoff between prefiltering fuel flow and varying the number of past measurements can be examined. The previous example, Fig. 8, was generated with a 20-s moving point mean applied to the fuel flow signal and the Kalman filter iterating every 20 s for all cases. With two measurements, by slowing the rate of iteration and allowing for more prefiltering, the performance can be improved. Three cases were examined to evaluate this tradeoff, and are presented in Table 1. Each case has a different frequency at which the Kalman filter is sampled, and the moving point mean time is adjusted accordingly to average over the entire gap between Kalman filter iterations. Case 1 samples the slowest but prefilters fuel flow longer, whereas Case 3 samples faster but prefilters less.

Table 1. Test cases to evaluate Kalman filter sample rate versus prefiltering.

Case	Time between iterations (moving point mean time, s)	Previous measurements, M
1	50	2
2	20	5
3	12.5	8

Figure 9 shows the results for each case, run for a total of 1000 s. For each case, the gain, G , and Kalman filter process covariance matrix, Q_k , were also tuned for best Kalman filter performance. The resulting trajectories all arrive at the minimum, however, the duration required to get there varies. Generally, Case 2 yielded the best overall performance when evaluated on steepest-descent trajectory and convergence to the minimum. Better Kalman filter performance could be achieved by both increasing the number of measurements and increasing the filtering time, however, it would then take longer to converge on the minimum. With the intention of flight-testing the algorithm, tests on the scale of 1000 s were desired to run multiple tests per flight, thus, the total time was not increased in order to improve performance.

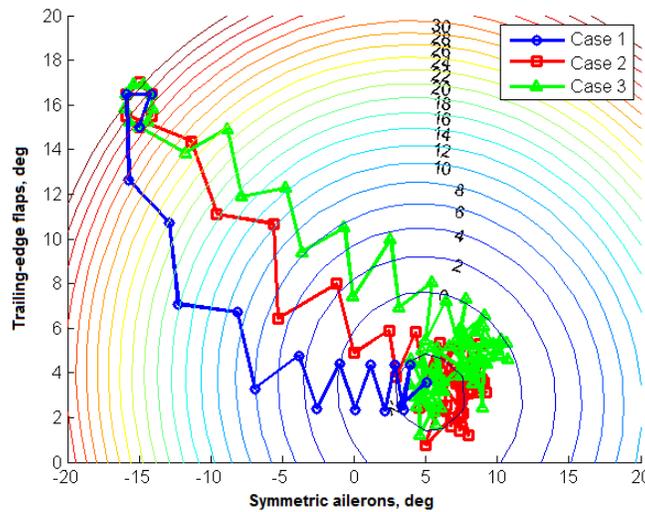


Figure 9. Three cases that exemplify the tradeoff between prefiltering versus increasing Kalman filter iteration frequency.

The value chosen for the noise covariance matrix, R , was determined based on the standard deviation of the fuel flow signal recorded during the performance function characterization flight as well as based on the length of the moving point mean applied to the fuel flow signal. The Kalman filter performance was then tuned by varying the diagonal components of the process covariance matrix, Q . Tuning was done by running multiple simulation runs at varying values of Q and evaluating the resulting performance based on the tuning metrics. The final tuned parameters of the system for the two-effector case based on simulation performance are given in Table 2. For the three-effector case, the same dwell time and the same Q and R matrices are used, however, M is increased and the gain is decreased. The general trend seen when increasing the number of effectors is the need to use more previous-data measurements and a lower gain.

Table 2. The final tuned parameter values for the two-effector case.

Parameter	Value
Gain	-105
M	5
Dwell time	20 s
R	$1.85^2 I$
Q	$1.98^2 I$

B. Simulation Results

The final tuned peak-seeking algorithm was evaluated in simulation by starting the system at four different initial conditions: one near the baseline trim of the airplane (Case A) and three cases offset from the baseline trim position with higher drag (Cases B through D). The three higher-drag initial condition cases (B through D) were chosen to evaluate the performance of the system at descending the performance function. Case A was chosen to evaluate the fine convergence and settling properties of the algorithm. It was also used as a test to determine if a detectable amount of fuel flow reduction relative to the baseline trim of the aircraft could be generated from the algorithm. The performance model developed above demonstrated a potential for -2.3 percent in the two-effector mode, but Case A was designed to see how close the algorithm could get to that point when starting from the nominal production trim.

The results from the two-effector case, optimizing trailing-edge flaps and symmetric ailerons, with leading-edge flaps fixed at nominal, are presented in Fig. 10. It can be seen that after the initial excitation circle is performed and an initial local gradient is estimated, the trajectories head toward the lower fuel flow with the sinusoidal persistent excitation. Each case converges on an area that is below the fuel flow for the baseline trim of the aircraft. The total duration for each case is 15 min.

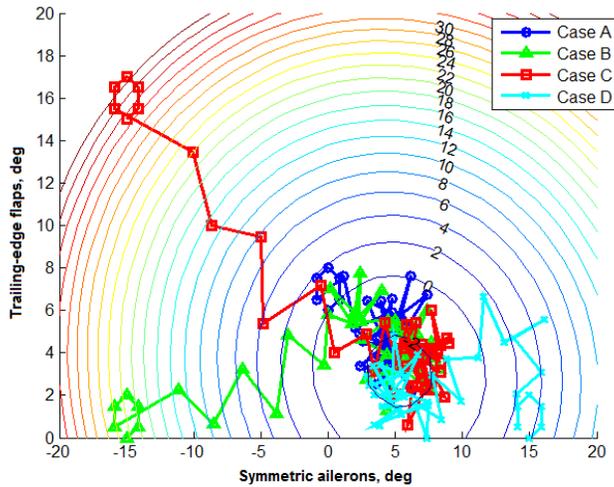


Figure 10. The peak-seeking algorithm trajectories from varying initial trim configurations.

The initial and final fuel flow values as a percentage from baseline for each case are shown in Table 3.

Table 3. Fuel flow reduction for Cases A through D.

Case	Initial trim	Initial fuel flow change, %	Final fuel flow change, %
A	Baseline	0	-2.2
B	High drag	+22	-2.2
C	High drag	+36	-2.1
D	Low drag	+3.5	-2.1

A time history of the fuel flow for Case B is shown in Fig. 11. The airplane starts initially from the baseline trim position. After 20 s the initial trim bias associated with Case B is inserted and the fuel flow increases. The initial excitation is then performed, followed by the algorithm engagement at 140 s. The fuel flow then rapidly decreases after only a few iterations and eventually converges around a fuel flow slightly below the baseline fuel flow for the airplane. While the unfiltered fuel flow signal varies by over ± 5 percent, the dwell time between iterations with a moving point mean reduces the effect of the noise. An even longer window average over several minutes shows a mean reduction of 2.2 percent from the baseline.

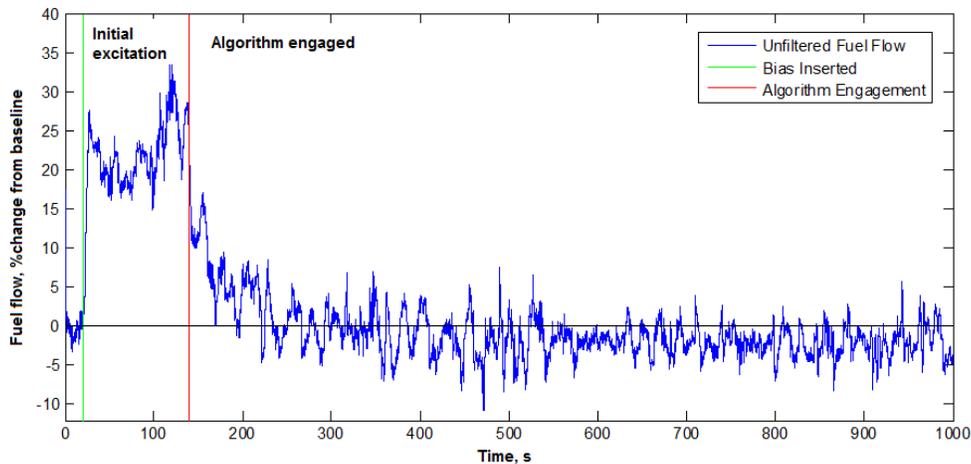


Figure 11. Time history of fuel flow for simulation run, Case B.

The results of the full three-effector case are shown in Fig. 12, Leading-edge flaps, trailing-edge flaps, and ailerons are all optimized simultaneously. Initially, the system performs a sphere of points to generate the local gradient. The system then descends toward the minimum and converges near the expected three-axis minimum fuel flow. The absolute minimum fuel flow expected based on the performance function model is -3.1 percent, however, the mean value that was achieved was -2.5 percent because the leading-edge flaps wandered about the minimum, between 5 and 9 deg, when the ideal position is 7 deg. This amount of reduction in fuel flow is close to the actual realized value found in flight which is detailed in a companion paper, Ref. 6.

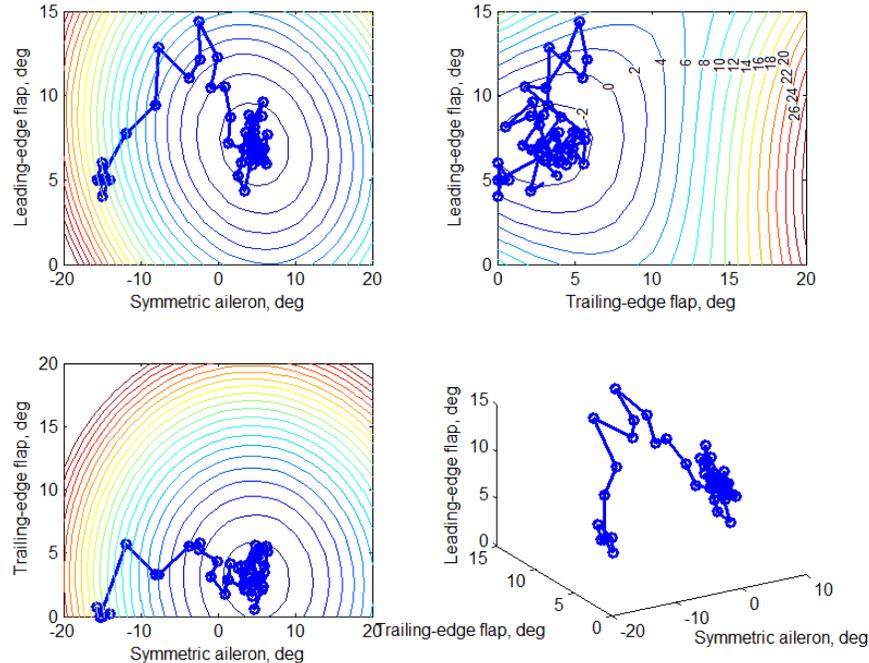


Figure 12. Three-control-effector optimization.

VII. Conclusions

The peak-seeking control algorithm presented in this paper was designed for real-time trim optimization for reduced fuel consumption. This algorithm has been tested and evaluated in a simulation environment using performance data collected in flight for an F/A-18 airplane (McDonnell Douglas, now The Boeing Company, Chicago, Illinois). With the algorithm engaged, the airplane was able to identify the local gradient of the performance function and descend to lower fuel flow conditions. Optimization of the leading-edge flaps, trailing-edge flaps, and symmetric aileron deflections showed potential for a 2.5-percent fuel flow reduction from the baseline trim configuration. The simulation tests were examined at a single flight condition; at other flight conditions fuel flow reduction may be greater. For other aircraft the amount of reduction may vary significantly, depending on how well the aircraft is optimized for its present weight configuration and flight condition. This algorithm could be integrated into the autopilot systems on some existing aircraft that have digital flight control systems and the ability to actuate multiple control surfaces symmetrically. Aircraft of the future can benefit from using this algorithm over existing methods of trim control by removing the need for the development of extensive lookup tables and by potentially providing an even more efficient trim configuration.

References

- ¹Collier, F., "Overview of NASA's Environmentally Responsible Aviation Project," 48th AIAA Aerospace Sciences Meeting, Jan. 2010.
- ²Federal Aviation Administration, "FAA Aerospace Forecast Fiscal Years 2013-2033," 2013, p.115.
- ³U.S. Environmental Protection Agency, "Inventory of U.S. Greenhouse Gas Emissions and Sinks: 1990-2011," April 2013, pp. 3-19 – 3-20.
- ⁴Ryan, J. J., and Speyer, J. L., "Peak-Seeking Control Using Gradient and Hessian Estimates," *American Control Conference*, Baltimore, Maryland, June 2010, pp. 611-616.

- ⁵Griffin, B. J., Brown, N. A., and Yoo, S. Y., "Intelligent Control for Drag Reduction on the X-48B Vehicle," AIAA-2011-6470, 2011.
- ⁶Brown, N. A., and Schaefer, J., "Peak-Seeking Optimization of Trim for Reduced Fuel Consumption: Flight-Test Results," AIAA Guidance, Navigation, and Control Conference, American Institute of Aeronautics and Astronautics, Reston, Virginia (submitted for publication).
- ⁷Phillips, P. W., and Smith, S. B., "AFTI/F-111 Mission Adaptive Wing (MAW) Automatic Flight Control System Modes Lift and Drag Characteristics," AFFTC-TR-89-03, May 1989.
- ⁸Boeing Advanced Systems Co., "AFTI/F-111 Mission Adaptive Wing Briefing to Industry," AFWAL-TR-88-3082, Oct. 1988.
- ⁹Renken, J., "Mission-Adaptive Wing Camber Control Systems for Transport Aircraft," AIAA-85-5006 1988.
- ¹⁰Thomson, K., and Schulze, E. T. "Delivering Fuel and Emissions Savings for the 777," Boeing Aeromagazine Quarter 3, 2009 pages 5-7
- ¹¹The Boeing Company, "Dynamic Adjustment of Wing Surfaces for Variable Camber," U.S. Patent Number 7,641,152, filed January 5, 2010.
- ¹²Zink, P. S., Love, M. H., and Youngren, H., "Drag Minimization Through the Use of Mission Adaptive Trailing Edge Flaps and Fuel State Control," AIAA-2004-4365, 2004.
- ¹³Gilyard, G., "In-Flight Transport Performance Optimization: An Experimental Flight Research Program and an Operational Scenario," NASA TM-97-206229, 1997.
- ¹⁴Gilyard, G., and Espana, M., "On the Use of Controls for Subsonic Transport Performance Improvement: Overview and Future Directions," NASA TM-4605, 1994.
- ¹⁵Gilyard, G. B., and Orme, J. S., "Performance Seeking Control: Program Overview and Future Directions," NASA TM-4531, 1992.
- ¹⁶Hanson, C., Ryan, J. J., "Peak-Seeking Optimization of Spanwise Lift Distribution for Wings in Formation Flight," AIAA-2012-4692, 2012.
- ¹⁷Hanson, C., "Capability Description for NASA's F/A-18 TN 853 as a Testbed for the Integrated Resilient Aircraft Control Project," DFRC-IRAC-CAP-002, 2009.
- ¹⁸Miller, C. J., "Nonlinear Dynamic Inversion Baseline Control Law: Architecture and Performance Predictions," AIAA-2011-6467, 2011.
- ¹⁹Ryan, J., and Speyer, J., "Systems and Methods For Peak-Seeking Control," U.S. Patent Number 8,447,443, May 21 2013.