

A methodology to assess the capability of engine designs to meet closed-loop performance and operability requirements

Alicia M. Zinnecker *

N&R Engineering, Parma Hts, OH, 44130, USA

and

Jeffrey T. Csank†

NASA Glenn Research Center, Cleveland, OH 44135, USA

Designing a closed-loop controller for an engine requires balancing trade-offs between performance and operability of the system. One such trade-off is the relationship between the 95% response time and minimum high-pressure compressor (HPC) surge margin (SM) attained during acceleration from idle to takeoff power. Assuming a controller has been designed to meet some specification on response time and minimum HPC SM for a mid-life (nominal) engine, there is no guarantee that these limits will not be violated as the engine ages, particularly as it reaches the end of its life. A characterization for the uncertainty in this closed-loop system due to aging is proposed that defines elliptical boundaries to estimate worst-case performance levels for a given control design point. The results of this characterization can be used to identify limiting design points that bound the possible controller designs yielding transient results that do not exceed specified limits in response time or minimum HPC SM. This characterization involves performing Monte Carlo simulation of the closed-loop system with controller constructed for a set of trial design points and developing curve fits to describe the size and orientation of each ellipse; a binary search procedure is then employed that uses these fits to identify the limiting design point. The method is demonstrated through application to a generic turbofan engine model in closed-loop with a simplified controller; it is found that the limit for which each controller was designed was exceeded by less than 4.76%. Extension of the characterization to another trade-off, that between the maximum high-pressure turbine (HPT) entrance temperature and minimum HPC SM, showed even better results: the maximum HPT temperature was estimated within 0.76%. Because of the accuracy in this estimation, this suggests another limit that may be taken into consideration during design and analysis. It also demonstrates the extension of the characterization to other attributes that contribute to the performance or operability of the engine. Metrics are proposed that, together, provide information on the shape of the trade-off between response time and minimum HPC SM, and how much each varies throughout the life cycle, at the limiting design points. These metrics also facilitate comparison of the expected transient behavior for multiple engine models.

Nomenclature

Variables

R_ψ	rotational transformation matrix from (x,y) to (\hat{x},\hat{y})
T_{40}	temperature at high-pressure turbine inlet
h	vector containing all engine health parameters
h_{eol}	engine health parameter at end-of-life condition

*Control Systems Engineer, alicia.m.zinnecker@nasa.gov, AIAA Member.

†Research Engineer, Intelligent Control and Autonomy Branch, jeffrey.t.csank@nasa.gov, AIAA Member.

$maxT40$	maximum HPT entrance temperature attained during transition from idle to maximum thrust ($^{\circ}R$)
$minSMa$	minimum HPC surge margin attained during transition from idle to maximum thrust (%)
$minSMd$	minimum HPC surge margin for which a controller's acceleration schedule has been designed (%)
r_x	half-length of the ellipse x -axis
r_{yb}	length of the bottom half of the ellipse y -axis
r_{yt}	length of the top half of the ellipse y -axis
tr	time for thrust to reach 95% of maximum, given demand from idle to maximum (sec)
\bar{x}^*	minimum allowable HPC surge margin in worst-case (horizontal minimum)
(x,y)	coordinates representing an ellipse centered at the origin, with axes coincident to fixed x - and y -axes
(\bar{x},\bar{y})	coordinates representing an ellipse centered at the design point, with axes rotated counter-clockwise ψ radians relative to the fixed x - and y -axes
(\hat{x},\hat{y})	coordinates representing an ellipse centered at the origin, with axes rotated counter-clockwise ψ radians relative to the fixed x - and y -axes
\bar{y}^*	maximum allowable response time (vertical maximum) in worst-case
ψ	tilt, in radians, of the ellipse x -axis relative to a fixed x -axis

Subscripts

b	points along the bottom half of the ellipse (in (x,y) coordinate system)
c	points at the center of the ellipse (minimum HPC surge margin or response time of nominal engine)
$max+$	point furthest 'above' the ellipse x -axis; point with the most positive y -coordinate in the (x,y) system
$min-$	point furthest 'below' the ellipse x -axis; point with the most negative y -coordinate in the (x,y) system
t	points along the top half of the ellipse (in (x,y) coordinate system)

Acronyms

C-MAPSS40k	Commercial Modular Aero-Propulsion System Simulation, 40k
HPC	High-Pressure Compressor
HPT	High-Pressure Turbine
LPC	Low-Pressure Compressor
RMS	Root-Mean-Square
SM	Surge Margin
TTECTrA	Tool for Turbine Engine Closed-loop Transient Analysis

I. Introduction

IN designing a controller, it is often of interest to estimate the robustness of the closed-loop system in the presence of uncertainties, such as input or measurement disturbances or unmodeled dynamics or changes in the plant, which are unaccounted for during control design. One particular example of a controller that should be robust to such unmodeled changes is an aircraft engine controller. Of interest in this effort is the contribution of engine age to uncertainty in the system: the closed-loop controller designed at a nominal operating point should provide the desired level of performance as the engine ages. Present engine designs, in part, utilize steady-state operating margins defined based on an assumed aging trend to take this uncertainty into account; dynamic performance is not considered.¹ The method proposed here characterized uncertainty due to aging using the results of dynamic system simulations to provide information on these margins and demonstrate how robust the system is to aging.

Previous work regarding robust control for uncertain systems has focused on plants represented by linear

models.²⁻⁴ Two control schemes often applied to linear models, \mathcal{H}_∞ control and μ synthesis, depend on the state-space matrices for development of a controller and to show stability or performance robustness. In either scheme, uncertainty must be expressed in a closed form, such as parametric or multiplicative, or through exogenous inputs to the plant. The latter representation was used in Refs. 5,6 in applying \mathcal{H}_∞ control to a piecewise-linear engine model to design gain scheduled multi-variable controllers. Many of the proposed extensions of these schemes to nonlinear models place stringent requirements on the systems in order for the controller to meet the design objectives.⁷⁻¹⁰ For example, the work in Ref. 7 was derived for discrete-time piecewise linear systems, while the effort described in Ref. 8 used nonlinear models affine in the control input; application of these results to an engine model would require several unreasonable assumptions to be made. Refs. 9,10 presented case studies comparing techniques for analyzing robustness in linear and nonlinear systems. In both studies, it was shown that applying linear robust control techniques and analysis to nonlinear systems greatly reduced the guaranteed range of stability. The case study in Ref. 10 additionally demonstrated a technique that used set theory to design a controller directly from a nonlinear system; this controller was more robust than one designed using a more conventional approach.

An alternative to these analytical approaches is to characterize the effects of uncertainty or disturbances using data collected through Monte Carlo simulations. For a highly-nonlinear model that takes uncertainty into account through inputs to the system, such a data-based approach precludes the need for a mathematical representation of how uncertainty enters and affects the model. Instead, uncertainty is characterized through a known mathematical expression defining the boundary of the Monte Carlo data.

When applied to a closed-loop engine model, the proposed data-based characterization forms the basis of an approach for designing a controller that guarantees specified performance and operability limits are met throughout the engine life cycle. The relationship between performance and operability represents a trade-off made during control design; here the time required for the thrust to reach 95% of its maximum value when accelerating from idle, tr , is considered a measure of performance and the minimum surge margin (SM) reached in the high-pressure compressor (HPC) during this acceleration, $minSMa$, serves as a quantification of operability. Together, these define the performance level of the closed-loop system for the given controller. A typical control design involves scheduling the gains of individual controllers constructed at specific points in the flight envelope; these operating points are defined by the environmental conditions and power level at which the engine is operating. In particular, it is not possible to design a controller that simultaneously allows for an arbitrarily fast response without causing the HPC to surge; instead, the controller must be designed for acceptable transient performance without violating a constraint on SM. A similar trade-off exists between the minimum HPC SM and the maximum HPT entrance temperature, $maxT40$, a limit that is of interest due to the reduction in turbine blade life when operated at high temperatures for long periods of time. Considering two controllers, one allowing for a fast response and one allowing sufficient HPC SM, an engine in closed-loop with the former, in addition to operating closer to surge, reaches a higher $maxT40$ than the latter. This risks more damage to the turbine blades during transients.

In the same way in which a controller designed to meet a specified performance level for one set of environmental conditions is not guaranteed to meet that same level at another set, the age of the engine affects closed-loop performance level. A controller designed to meet particular performance and operability limits in a new engine is not guaranteed to meet those same limits in an end-of-life engine, or in an engine operating under worst-case aging. Given a controller designed for a nominal engine (the design point), both trade-offs are negatively impacted as the engine ages. An older engine operates closer to surge (the minimum HPC SM reached during acceleration decreases) and the structural limits (the maximum HPT temperature increases). At the given design point, response time increases only slightly, because the controller was designed to achieve certain transient characteristics in the system. This suggests that, if it were possible to estimate the worst-case performance level at an arbitrary design point, a method could be developed which selects the design point at which a limit on response time or minimum HPC SM is guaranteed to be met throughout the life-cycle of the engine.

The uncertainty characterization proposed here is one such method for describing the relationship between tr and $minSMa$ for random aging in the engine. Application of this approach yields limiting design points that define bounds on the admissible controllers that are expected to not exceed a design limit on response time or minimum HPC SM. (The characterization can be extended to other representations of performance or operability, demonstrated through application to the trade-off between maximum high-pressure turbine (HPT) entrance temperature and minimum HPC SM, which provides additional information to consider during design or analysis of the system). Metrics are proposed that describe the trade-off between minimum

HPC SM and response time for the nominal engine and to capture how aging affects this relationship at specific design points. Each metric represents some measure of distance between the nominal design point and design limits and may be used to compare engine models.

The approach for characterizing uncertainty is detailed in Section II, beginning with a discussion of the motivation for why such a characterization was of interest. Evaluation of this approach through its application to a closed-loop turbofan engine model is presented in Section III, where the approach is extended to show how the maximum HPT entrance temperature can also be predicted. The results are summarized in Section IV with a discussion of future development and implementation of the proposed uncertainty approach.

II. Description of approach

At a fixed set of environmental conditions, an aircraft engine controller design involves selecting gains for the setpoint controller to meet tracking requirements in the presence of small disturbances, then designing acceleration and deceleration schedules to, among other considerations, protect against surge in the high- and low-pressure compressors for large transients. The acceleration schedule, in particular, has a noticeable effect on the performance level of the closed-loop system: designing the acceleration schedule to enforce a desired minimum HPC SM, $minSMd$ limits how quickly the engine can respond to a change in power setting; this $minSMd$ defines the design point for a given controller. For a set of trial design points, selected uniformly on the admissible range of $minSMd$ for a given engine model, the uncertainty characterization requires information about how engine aging affects the trade-offs between response time and minimum HPC SM, and between maximum HPT temperature and minimum HPC SM. This is facilitated through the use of an engine model with health parameter inputs that modify the flow rate and efficiency interpolated from the compressor and turbine maps during simulation. Each health parameter takes a value between 0, for a new engine, and h_{eol} , for an end-of-life engine and is uniformly sampled on the range 0 to h_{eol} to obtain the set of health parameter vectors, h , used for the Monte Carlo simulations. This guarantees that, when normalized with respect to the magnitude of the end-of-life engine health parameter vector, $||h|| \leq 1$. Additionally, data are collected for the nominal (mid-life), new, and end-of-life engines, providing relatively consistent endpoints of the performance level. Each simulation captures the engine response when commanded to accelerate from idle to takeoff; selection of the limiting design point considers a limit on tr or $minSMa$, but not both at the same time. (The proposed approach may be extended to instead consider both limits simultaneously, but the purpose of the present effort was to identify those points bounding designs that do not violate one of the two limits in the worst-case.)

The general approach involves three main steps:

1. Characterize the uncertainty by constructing elliptical regions that encompass the data collected at each design point.
2. Construct curve fits relating the design point to parameters describing the size and orientation of the ellipse.
3. Implement a binary search procedure that uses the curve fits to determine a design point meeting either a worst-case minimum HPC SM, \bar{x}^* , or a worst-case maximum response time, \bar{y}^* .

The approach is then evaluated through Monte Carlo simulations performed at each limiting design point identified in Step 3 above, followed by construction of the boundaries approximating this data. Comparison between the actual and estimated elliptical regions is used to determine how well the design limits are predicted.

A. Controller design and data collection

The uncertainty characterization relies on data from two sets of simulations: one set provides $minSMa$, tr , and $maxT40$ at three distinct points in the engine life-cycle (new, mid-life, and end-of-life) and the other provides the same information for randomly-aged engines (the Monte Carlo simulations). These simulations are performed at each trial design point for a given engine model; each design point is associated with a set of controller specifications. These define the bandwidth and phase margin for the setpoint controller, as well as the minimum compressor SMs for which the acceleration and deceleration schedules are constructed. (Ideally, these will be the minimum HPC SM attained in the system during transients; that is, for the HPC,

$minSMd = minSMa$ for a nominal engine.) For simplicity, the same setpoint controller and deceleration limiter are used for each trial design point for a given engine model. Since only the acceleration schedule differs between controllers, the design points are identified concisely by the absolute minimum HPC SM at which the nominal engine should operate when subject to the schedule (that is, $minSMd$). The trial design points are chosen uniformly on an admissible range of $minSMd$ for the given engine. The range should provide enough information to capture the trade-off between tr and $minSMa$ for design points that do not exceed the limits \bar{x}^* and \bar{y}^* .

B. Mathematical background

To maintain flexibility in the uncertainty characterization, boundaries defined by the union of two half-ellipses, rotated and shifted to the nominal performance level (\bar{x}_c, \bar{y}_c) , are considered. These define the region in the $(minSMa, tr)$ plane where the Monte Carlo simulation data for a specific design point lay. Each ellipse is defined by four parameters: the lengths of the ellipse x -axis, r_x , the upper-half ellipse y -axis, r_{yt} , and the lower-half ellipse y -axis, r_{yb} , and the tilt of the ellipse x -axis in relation to the horizontal axis, ψ . To define these parameters from the Monte Carlo data, three coordinate systems are defined, which are related through linear and rotational transformations; these are illustrated in Figure 1. The ‘fixed coordinate system,’ designated by unbarred variables, is the $(minSMa, tr)$ plane; for simplicity, x and y are used for these coordinates, respectively. The ‘ellipse coordinate system,’ indicated by barred variables, defines the location and orientation of the elliptical boundary with respect to the fixed system. An intermediate system, denoted by hatted variables, defines the linear translation of the ellipse coordinate system from (\bar{x}_c, \bar{y}_c) to the origin. The blue solid, green dotted, and red dashed curves in the figure represent an ellipse in each coordinate system, respectively. Indicated on the ellipse in the fixed coordinate system are r_x , r_{yt} , r_{yb} , and ψ , which are calculated from the simulation data collected at each design point and defines the uncertainty at that point. When the Monte Carlo simulation data lay in a more circular region, the axis lengths would be nearly equal, while an aging trend more concentrated along a straight line may result in r_x being much larger than r_{yt} and r_{yb} (or vice versa).

The performance levels of the new, mid-life, and end-of-life engines, the first set of simulation data described in Section II.A, are used to determine r_x and ψ , where the nominal performance level is determined from the mid-life engine. The ellipse x -axis is defined by the line through (\bar{x}_c, \bar{y}_c) that is parallel to that connecting the new and end-of-life performance levels. The length r_x is defined as the largest distance between the projection of these points to the ellipse x -axis and the center of the ellipse. The angle through which the axes are rotated, ψ , is the arctangent of the slope of the ellipse x -axis.

To calculate the lengths of the two ellipse y -axes, the ellipse equation is invoked; this requires transfor-

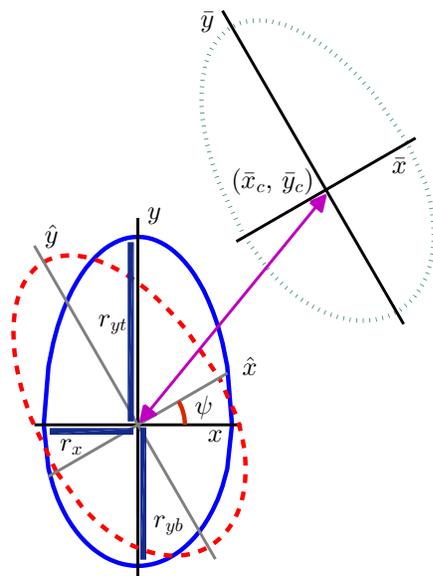


Figure 1. Illustration of the ellipse parameters r_x , r_{yt} , r_{yb} , and ψ , the coordinate systems (x, y) , (\hat{x}, \hat{y}) , and (\bar{x}, \bar{y}) , and the translational transformation (x_c, y_c) used to define the elliptical boundaries for each region of simulation data.

mation of the Monte Carlo simulation data from the ellipse coordinate system, which is centered at (\bar{x}_c, \bar{y}_c) and rotated counter-clockwise ψ radians from the horizontal axis (the ‘ellipse coordinate system,’ indicated by barred variables), to the fixed coordinate system. After a linear translation of the Monte Carlo simulation data to the intermediate system, the following rotation matrices define the relationship between the intermediate and fixed systems:

$$R_\psi = \begin{bmatrix} \cos \psi & -\sin \psi \\ \sin \psi & \cos \psi \end{bmatrix} \quad R_\psi^{-1} = \begin{bmatrix} \cos \psi & \sin \psi \\ -\sin \psi & \cos \psi \end{bmatrix} \quad (1)$$

where R_ψ transforms (x, y) to (\hat{x}, \hat{y}) , and R_ψ^{-1} performs the inverse transformation. To define r_{yt} and r_{yb} , the data points furthest ‘above’ and furthest ‘below’ the ellipse x -axis are identified. This is most easily done in the fixed coordinate system, where these are the points with the most positive y -coordinate, (x_{max+}, y_{max+}) , and with the most negative y -coordinate, (x_{min-}, y_{min-}) , respectively. Given r_x and these coordinates, the two ellipse y -axis lengths are found using Equation (2).

$$r_{yt} = \frac{(y_{max+})r_x}{\sqrt{r_x^2 - (x_{max+})^2}} \quad (2)$$

$$r_{yb} = -\frac{(y_{min-})r_x}{\sqrt{r_x^2 - (x_{min-})^2}}$$

Let the parameter $x \in [-r_x, r_x]$ define the range on which each half-ellipse is defined. Transforming the curve to the ellipse coordinate system through application of R_ψ and a linear translation, the ellipse is defined as the union of the parametric equations:

$$\begin{aligned} \text{top of ellipse : } & \begin{cases} \bar{x}_t = x \cos \psi - \frac{r_{yt}}{r_x} \sin \psi \sqrt{r_x^2 - x^2} + \bar{x}_c \\ \bar{y}_t = x \sin \psi + \frac{r_{yt}}{r_x} \cos \psi \sqrt{r_x^2 - x^2} + \bar{y}_c \end{cases} \\ \text{bottom of ellipse : } & \begin{cases} \bar{x}_b = x \cos \psi + \frac{r_{yb}}{r_x} \sin \psi \sqrt{r_x^2 - x^2} + \bar{x}_c \\ \bar{y}_b = x \sin \psi - \frac{r_{yb}}{r_x} \cos \psi \sqrt{r_x^2 - x^2} + \bar{y}_c \end{cases} \end{aligned} \quad (3)$$

from which the extrema can be found. For this uncertainty characterization, the vertical maximum (worst-case response time) and horizontal minimum (worst-case minimum HPC SM) are of particular interest. Because $\psi \in [-\frac{\pi}{2}, \frac{\pi}{2}]$, the vertical maximum is always on the top half of the ellipse and its coordinates are found using the first pair of expressions in Equation (3). In contrast, the sign on ψ determines on which half of the ellipse the horizontal minimum is located: if $\psi > 0$, the minimum occurs on the top half, and the reverse is true when $\psi < 0$. The coordinates of each extrema are listed in Equation (4).

$$\begin{aligned} \text{vertical maximum : } & \left(\frac{(r_x^2 - r_{yt}^2) \sin \psi \cos \psi}{\sqrt{r_x^2 \sin^2 \psi + r_{yt}^2 \cos^2 \psi}} + \bar{x}_c, \sqrt{r_x^2 \sin^2 \psi + r_{yt}^2 \cos^2 \psi} + \bar{y}_c \right) \\ \text{horizontal minimum : } & \begin{cases} \left(-\sqrt{r_x^2 \cos^2 \psi + r_{yt}^2 \sin^2 \psi} + \bar{x}_c, -\frac{(r_x^2 - r_{yt}^2) \sin \psi \cos \psi}{\sqrt{r_x^2 \cos^2 \psi + r_{yt}^2 \sin^2 \psi}} + \bar{y}_c \right) & \text{when } \psi > 0 \\ \left(-\sqrt{r_x^2 \cos^2 \psi + r_{yb}^2 \sin^2 \psi} + \bar{x}_c, -\frac{(r_x^2 - r_{yb}^2) \sin \psi \cos \psi}{\sqrt{r_x^2 \cos^2 \psi + r_{yb}^2 \sin^2 \psi}} + \bar{y}_c \right) & \text{when } \psi < 0 \end{cases} \end{aligned} \quad (4)$$

C. Defining the regions for an engine model

For a given engine model, simulation results are collected at a series of trial design points, uniformly selected on an admissible range of $minSMd$ for the given model. Repeated application of the derivation in Section II.B defines the regions on the plane in which each set of data is located. Each trial design point is then associated with seven values: the HPC SM limit for which the acceleration schedule was designed, $minSMd$, the location of the ellipse center, $(minSMa, tr)$, and four parameters related to the size and orientation of the ellipse, r_x , r_{yt} , r_{yb} , and ψ .

To extend this characterization to an arbitrary design point, a least-squares fitting algorithm is applied to construct curve fits relating the minimum HPC SM for which the controller was designed to the size and orientation of the ellipse defining the boundary for worst-case performance level at the design point. This involves constructing polynomials of various orders and selecting the one that best fits the data in a root-mean-square (RMS) sense. Two single-variable polynomials, $minSMa(minSMd)$ and $tr(minSMa)$,

approximate the center of the ellipse given a design point $minSMd$. The ellipse is defined through four multi-variable curves in $minSMa$ and tr . These six polynomials provide an estimation of the uncertainty in performance level at an arbitrary design point $minSMd$, from which the worst-case performance level can be estimated through the use of Equation (4).

D. Finding the design points

The ultimate goal of this uncertainty analysis is to provide a method for identifying the design points at which specific limits in $minSMa$ or tr , denoted \bar{x}^* and \bar{y}^* , respectively, are met. This will allow for definition of a range of possible design points for which at least one of the limits will not be exceeded in the worst-case. The curve fits discussed in Section II.C, together with the coordinates of the extrema in Equation (4), are incorporated as part of a binary search procedure to approximate these limiting design points. A range of $minSMd$ is determined from the simulation results to define the initial search range; on each iteration, the extremum is calculated using the ellipse approximated by the curve fits. The search range is adjusted on each iteration, and the procedure terminates when the extremum meets the limit, within a specified tolerance.

This approach is evaluated by designing a controller at each limiting design point and performing Monte Carlo simulations of the closed-loop engine models. The accuracy of the characterization is determined through comparison of the ellipses fit to this data to those estimated by the curve fits. The approach can be extended to evaluate the approximation of other limiting values in the system, such as the worst-case maximum HPT entrance temperature, minimum low-pressure compressor (LPC) SM, or minimum fuel-to-air ratio. Application to analysis of maximum HPT entrance temperature is presented in Section III.F as demonstration of this ability to estimate other limits in the system that may be of interest when selecting design points or analyzing the system.

III. Application to engine models

The uncertainty characterization and analysis approach described in Section II was applied to the engine model from the Commercial Modular Aero-Propulsion System Simulation, 40k (C-MAPSS40k) integrated with a controller designed using the Tool for Turbine Engine Closed-loop Transient Analysis (TTECTrA). Not only did this provide verification of the proposed approach, it also led to the definition of three metrics that describe the trade-off between tr and $minSMa$ and the robustness of the system to changes due to aging. These metrics provide a measure of comparing models, specifically in terms of how sensitive response time is to changes in the acceleration schedule (which determines $minSMa$), how much performance level is affected by degradation due to engine aging, and how tight the range of allowable control designs (those that do not exceed limits \bar{x}^* or \bar{y}^*) is.

A. Engine models

Each of the three engine models considered in this effort were derived from C-MAPSS40k, a zero-dimensional model of a twin-spool, high-bypass turbofan engine. The model is capable of being simulated at various stages of the engine life through the use of 10 health parameters, which define the degradation of efficiency or flow rate in the compressors and turbines.¹¹ The nominal engine, for which each controller was designed, was assumed to have a health parameter vector representing a mid-life engine. The first model, referred to as model 1 in the following discussion, is the baseline C-MAPSS40k model. Two additional models were obtained by adjusting the scalars used to size the values interpolated from the compressor and turbine maps that define the engine model. (These scalars affected the shaft speed, pressure ratio, flow rate, and efficiency relationship for each of the three compressors and two turbines in the engine.¹¹) Note that the choice of scalars for these two models was made to obtain models with response characteristics that differed from those of model 1 and serve mainly as design examples. Model 2 was created by adjusting scalars in every component, as well as the inertias of the two shafts, from those used in model 1. The efficiency scaler was increased for each compressor and turbine. Additionally, for both the high-pressure compressor and turbine, the flow rate scaler was decreased and the corrected shaft speed scaler was increased. For model 3, the scalars affecting the flow rate for the high- and low-pressure turbines were reduced from those used for model 1; no other modifications were made. The changes made to models 2 and 3 produced engines with significantly different dynamics from model 1, as shown in the results from application of the uncertainty characterization in Section III.D.

Table 1. Specifications provided to TTECTrA for designing controllers for each engine model.

model	bandwidth (Hz)	phase margin (deg)	minimum HPC SM (%)	minimum LPC SM (%)
1	1.5	45	5 to 18	15
2	1.5	45	4 to 13	15
3	1.5	45	2 to 11	15

B. Controller design

The baseline C-MAPSS40k controller is typical for a 40,000 lb_f thrust class high-bypass turbofan engine,¹² but was replaced for this effort by a controller designed using TTECTrA, which offers a simplified architecture.^{1,13} In particular, this controller contains three main components: a setpoint controller, an acceleration schedule, and a deceleration schedule. (The C-MAPSS40k controller contains these components in addition to several other limiters, which serve mainly to protect against structural damage and have limited effect on the transient behavior of the system.) The setpoint controller is gain-scheduled to enforce a particular bandwidth and phase margin on the response to small transients, while the limiters, in part, protect against surge in the high- and low-pressure compressors, respectively. TTECTrA designs the acceleration schedule, which specifies the maximum allowable acceleration of the core shaft at a given rotational speed, using criteria on the minimum HPC SM. The deceleration schedule is defined as the lower limit on the ratio between fuel flow and static pressure in the burner required to preserve a minimum SM in the LPC. Each controller is designed for a fixed set of environmental conditions; here, the sea level static standard-day conditions (0 ft. altitude, 0 Mach, 59°F) are assumed.

For simplicity, because the acceleration schedule has the most notable effect on performance level of the closed-loop system, the setpoint controller and deceleration limiter are the same for each controller constructed for a given model; the design specifications are listed in Table 1. The range of *minSMd* for which the acceleration schedules were calculated was determined on a model-specific basis to ensure that the limiting design points fell within the range for which the curve fits were constructed; the trial design points were chosen uniformly (in this case, in increments of 1%) on this range.

C. Simulation details

For each engine model, the two sets of simulations described in Section II.A were run. To obtain results for the new, mid-life, and end-of-life engines, each element of the health parameter vector was assigned a value based on the engine age; for example, $h = 0$ for a new engine and each element took on its largest value (h_{eol}) for the end-of-life engine. The particular value of h_{eol} for each parameter is unique to C-MAPSS40k and not included here for compactness. (It was assumed that each health parameter took a value on the same range in all three models, so that the same set of health parameter vectors was used in the Monte Carlo simulations for each model.)

To perform the Monte Carlo simulations, a set of 1000 health parameter vectors was constructed by independently, uniformly sampling each element on its range (0 to h_{eol}) to represent random aging in the engine. Unlike the piecewise linear aging trend typically assumed in C-MAPSS40k, where the elements of the health parameter vector all change proportionally to the engine age, this sampling process does not couple the elements. This may lead to the use of an unrealized health parameter vector, wherein, for example, a component may have a ‘fully degraded’ flow rate but a ‘new’ efficiency. This may be of interest to consider in future applications of the approach.

D. Fitting and binary search results

Results from simulation of the new, mid-life, and end-of-life engines, and from Monte Carlo simulations, were used in the calculations outlined in Section II.B to define the boundaries approximating the worst-case performance level of the system. Figures 2 to 4 show the data and calculated values for models 1, 2, and 3, respectively; the plots are color-coded by trial design point. In each figure, the Monte Carlo simulation results are shown as dots; the new, mid-life, and end-of-life performance levels are indicated by six-point stars, five-point stars, and triangles, respectively; and the ellipse is plotted as a solid line, with axes (dashed lines) and extrema (squares) overlaid. For each model, the boundaries were estimated relatively accurately: out of 14,000 data points collected for each model, 99.71% were contained within the respective regions estimated for model 1, 99.86% for model 2, and 99.59% for model 3. This suggested that the approximated regions for the limiting design points should provide reasonable estimates of the worst-case simulation results.

Two predominant observations made from Figures 2 to 4 regard the illustrated trade-off between response

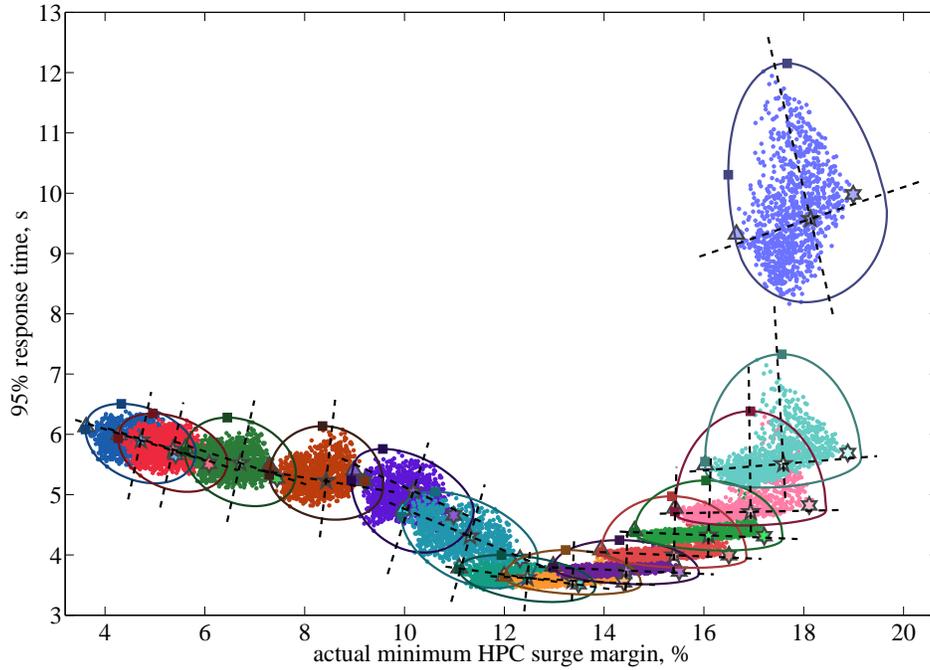


Figure 2. Plot showing data from the Monte Carlo simulations of model 1, color-coded by design point, overlaid with data for new, mid-life (nominal), and end-of-life engines (six-point stars, five-point stars, and triangles, respectively) and the ellipses fit to the data (solid lines). Ellipse axes are indicated by dashed lines, and extrema by boxes.

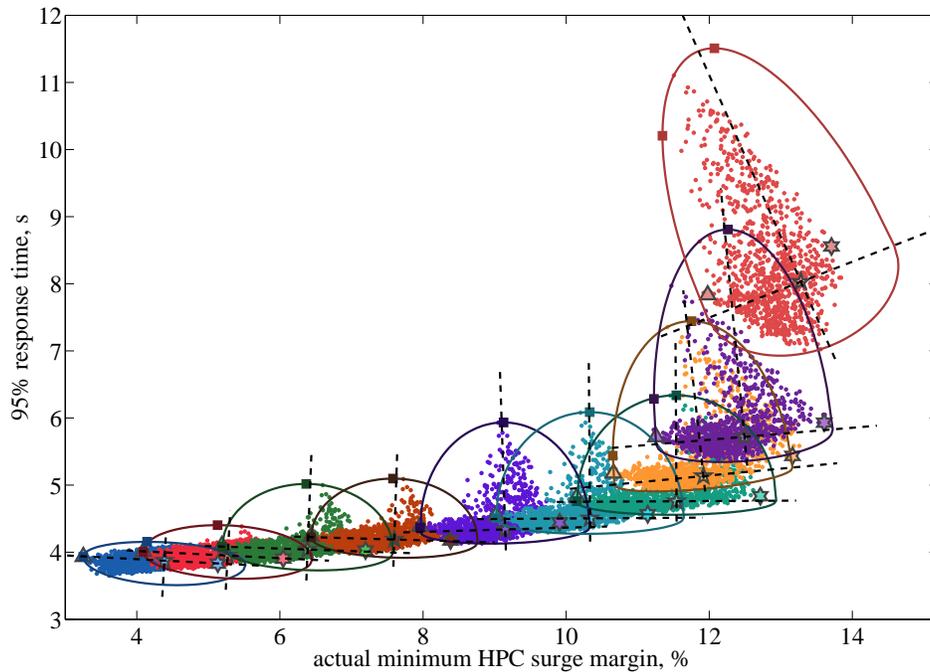


Figure 3. Plot showing data from the Monte Carlo simulations of model 2, color-coded by design point, overlaid with data for new, mid-life (nominal), and end-of-life engines (six-point stars, five-point stars, and triangles, respectively) and the ellipses fit to the data (solid lines). Ellipse axes are indicated by dashed lines, and extrema by boxes.

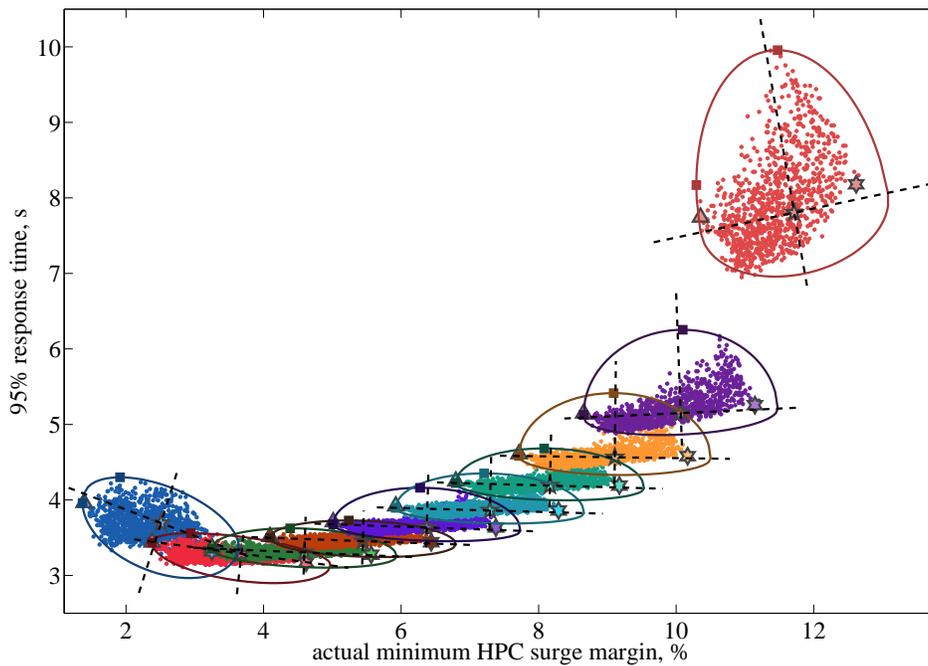


Figure 4. Plot showing data from the Monte Carlo simulations of model 3, color-coded by design point, overlaid with data for new, mid-life (nominal), and end-of-life engines (six-point stars, five-point stars, and triangles, respectively) and the ellipses fit to the data (solid lines). Ellipse axes are indicated by dashed lines, and extrema by boxes.

time and minimum HPC SM and the presence of uncertainty in the systems. The performance levels of each engine were compared by considering how response time and minimum HPC SM are related in each model. Generally, this relationship is expected to be monotonically increasing, so that, as the minimum HPC SM increases, the engine response slows down. This was only strictly true for model 2 (Figure 3), where the rate in which response time increased with respect to minimum HPC SM was relatively constant until the minimum HPC SM exceeds 11%. At this point, response time rapidly increased with minimum SM. The other two models deviated from this expected behavior due to integrator windup when operated at low minimum HPC SMs. Controllers designed using TTECTrA contain an anti-windup scheme to mitigate the integrator windup that occurs when one of the limiters is active. The gains of the anti-windup compensator for the setpoint controller are iteratively tuned to meet a pre-specified transient requirement, while those for the compensator in the acceleration limiter are fixed by TTECTrA. Although generally successful in reducing windup in the system, when $minSMd$ (and, consequently, $minSMa$) were decreased below a specific value, the compensators no longer functioned as expected, the result of which manifested itself through an increase in response time. It is important to note that, because this effect was seen in the data used to construct the curve fits, and because it appeared to occur in a predictable manner (for a given model), this had no effect on the results presented here. Model 2 was the only model for which this effect was not observed; for model 3 (Figure 4), this occurred for minimum SM less than 4%, while in model 1 (Figure 2) the effect was seen at much higher minimum SM (around 12%). As $minSMa$ increased above this value, the trade-off curves for both models followed a similar trend to model 2, appearing nearly linear until the minimum HPC SM increased above some value (17% for model 1, 10% for model 3), where the rate of change in response time rapidly increased. The choice of scalers in the three models appeared to affect how quickly response time increased with minimum HPC SM, and how likely windup was likely to occur at smaller minimum HPC SM.

The general effects of aging on the performance level were similar between the three engine models. Three distinct segments, related to the shape of the trade-off curve, were observed. At those design points where windup was present, aging caused the data to be distributed more evenly about the center, producing regions that were closer to circles than in other segments; this particular segment was not seen in the results for model 2. For those points on the more linear segment of the trade-off curve, where windup was not present, the ellipses became elongated along the x -axis, more closely approximating the aging trend suggested by considering the new, mid-life, and end-of-life engine performance levels exclusive of the Monte Carlo results. As the sensitivity of the trade-off increased, the data points became more spread out and formed a ‘peak’

that was not seen in the other two segments, making the ellipse elongated along the y -axis.*

The relationship between response time and minimum HPC SM, and the size and orientation of the elliptical regions for each design point, was approximated through the curve fits constructed using the data in Figures 2 to 4. For each model, the six curve fits relating the nominal relationship between $minSMd$, $minSMa$, tr , and the four ellipse parameters were constructed using the approach described in Section II.C. For the two single-variable polynomials, which represent the performance and operability trade-off, the fitting algorithm tested functions up to order 12 for each model. Multi-variable polynomials up to order 6 in each independent variable, with cross-terms up to order 3 in each variable (that is, $minSMa^3 \times tr^3$) were considered to estimate the ellipse parameters. Table 2 lists the orders of the best fit polynomials for each case; the orders listed for the multi-variable functions are for $minSMa$, tr , and cross-terms $minSMa^i \times tr^j$, respectively. Also listed are the RMS errors for each fit, which provide an indication of how closely the data used to construct the curve is matched by the fit.

The single-variable fits are compared to the data used to generate them in Figure 5, where the top row of plots shows $minSMa(minSMd)$ and the bottom row shows $tr(minSMa)$. Although relatively accurate, there was some disagreement between the actual response time and that predicted by $tr(minSMa)$ for models 1 and 2. This manifested itself in the RMS errors for these fits, which were 10 to 11 orders of magnitude larger than for model 3. Contributions to these relatively large RMS errors were due to mismatches between the actual and predicted tr for minimum HPC SM between 11% and 17% for model 1 and between 11% and 12% for model 2. Additionally, the fits for $minSMa(minSMd)$ for these models, particularly model 2, also show some inaccuracies. As a result, the location of design points on these ranges were predicted inaccurately, leading to under- or overestimation of the vertical maximum or horizontal minimum and, consequently, identification of a limiting design point that did not meet the limit as well as predicted. This will be addressed in Section III.E, with focus on the results for model 1.

Figures comparing the multi-variable fits to the data are not shown here for compactness, since agreement was extremely close as indicated through the small RMS errors listed in Table 2. These results suggested that inaccuracy in the predicted limiting design point could be attributed to errors from the single-variable fits that propagated to these multi-variable fits. This was not strictly true, however, due to the shape of the surface defined by the multi-variable function: although the data were fit well, the gradient of the surface local to the trial design points was large. In some cases, this caused small errors in $minSMa$ or tr (from the single-variable fits) to be magnified by the multi-variable fit; Figure 6 shows this for the fit $r_{yt}(minSMa, tr)$ for model 1, which is representative of each multi-variable function fit to data for each model. The effect of this large gradient was seen most significantly through large errors in r_{yt} or r_{yb} resulting from much smaller errors in $minSMa$ or tr , as will be addressed in Section III.E.

To identify the limiting design points that represent bounds for each model on the control designs that do not exceed either a minimum HPC SM or maximum response time limit in the worst-case, it was necessary to first define these limits. The latter limit was taken from Section 33.73 of Title 14 of the Code of Federal Regulation,¹⁴ which stated that acceleration from flight idle to 95% of the maximum thrust must occur within 5 seconds from application of the command; that is, the 95% response time must not exceed $\bar{y}^* = 5$ seconds. The limit on minimum HPC SM was selected at $\bar{x}^* = 7\%$, to allow the system to tolerate decreases in SM due to other sources of uncertainty not considered by this characterization. The vertical maximum

*This artifact motivated the use of two ellipse y -axis lengths instead of a single one because using a single length produced either an overestimate of the data in the bottom half of the ellipse, or an underestimate of the top half.

Table 2. Orders and RMS errors for each best-fit curve constructed for the C-MAPSS40k engine model. For the multi-variable polynomials (r_x , r_{yt} , r_{yb} , and ψ), a list is given which specifies the orders of the terms $minSMa$, tr , and $minSMa^i \times tr^j$, respectively.

Curve	Model 1		Model 2		Model 3	
	order	RMS error	order	RMS error	order	RMS error
$minSMd(minSMa)$	8	0.04563	8	0.0558	9	$3.3519 \cdot 10^{-12}$
$tr(minSMa)$	8	0.1144	8	0.03174	9	$5.6018 \cdot 10^{-11}$
$r_x(minSMa, tr)$	5, 2, (2, 3)	$3.7057 \cdot 10^{-13}$	5, 3, (1, 1)	$3.0790 \cdot 10^{-14}$	4, 2, (1, 3)	$2.7357 \cdot 10^{-15}$
$r_{yt}(minSMa, tr)$	5, 2, (2, 3)	$3.7057 \cdot 10^{-13}$	5, 1, (1, 3)	$6.2547 \cdot 10^{-14}$	4, 2, (1, 3)	$1.8259 \cdot 10^{-13}$
$r_{yb}(minSMa, tr)$	5, 4, (2, 2)	$3.7057 \cdot 10^{-13}$	5, 1, (1, 3)	$3.9480 \cdot 10^{-14}$	4, 2, (1, 3)	$1.585 \cdot 10^{-14}$
$\psi(minSMa, tr)$	5, 2, (2, 3)	$3.7057 \cdot 10^{-13}$	5, 1, (1, 3)	$3.2888 \cdot 10^{-14}$	4, 2, (1, 3)	$9.5555 \cdot 10^{-15}$

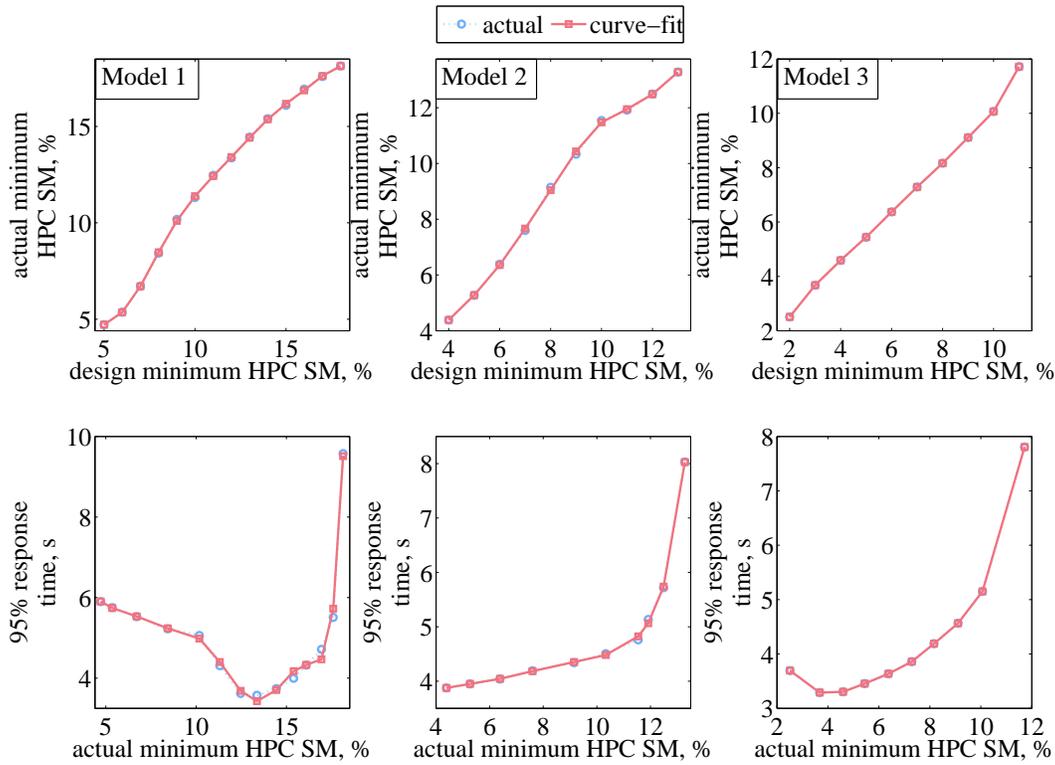


Figure 5. Comparison of the accuracy of the fits for $minSMa(minSMd)$ (top row) and $tr(minSMa)$ (bottom row) for models 1, 2, and 3, respectively columnwise.

and horizontal minimum of the elliptical regions for each trial design point were calculated for each model (Equation (4)) to identify the ranges of $minSMd$ on which the searches were to be performed. For model 1, the design point meeting \bar{x}^* was identified through a search of $minSMd$ between 7% and 8%. Because of the significant presence of windup at low minimum HPC SM for this model, two design points were identified at which the limits on \bar{y}^* were expected to be met: one between 10% and 11% and one between 14% and 15%. For model 2, the searches were performed for $minSMd$ between 7% and 8% and between 5% and 6%, respectively, and for model 3 both design points were on the range $minSMd = 8\%$ to $minSMd = 9\%$.

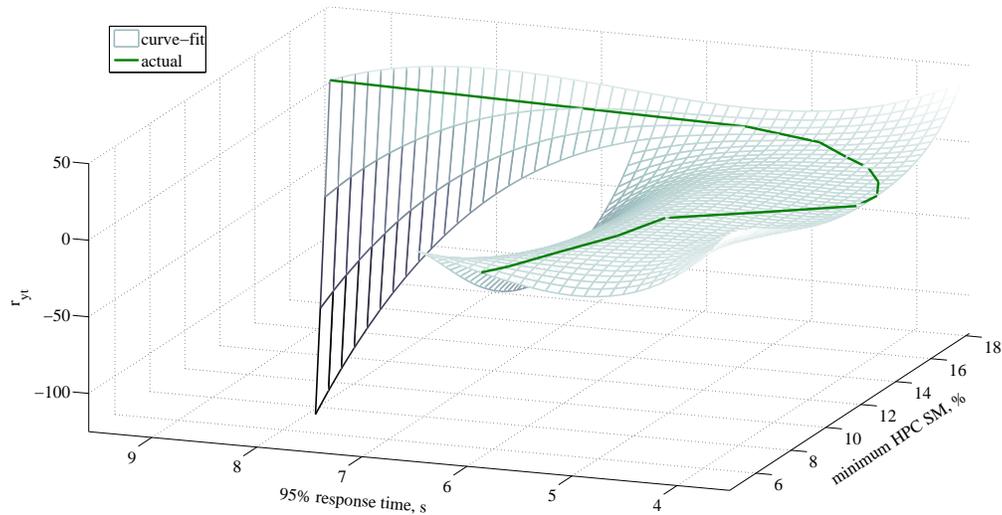


Figure 6. Illustration of the surface $r_{yt}(minSMa, tr)$ for model 1 with the data used to generate the fit overlaid to show how steep the gradient of the fit is for points just off the set of trial design points. This is representative of the results for all four multi-variable fits for each of the three models under consideration.

Table 3. Estimated nominal (fifth and sixth columns) and worst-case (third column) performance levels for the limiting design points for which $\bar{x}^* = 7\%$ or $\bar{y}^* = 5$ seconds are predicted to be met in each of the three models.

Model	Objective	Predicted limit	$minSMd$ (%)	$minSMa$ (%)	tr (sec)	search tolerance	# iterations
1	$\bar{x}^* = 7$	7.0032	7.8047	8.1207	5.2622	0.001	7
	$\bar{y}^* = 5$	4.9933	10.0625	11.4543	4.3079	0.004	4
		5.0074	14.7813	16.0121	4.3215	0.004	5
2	$\bar{x}^* = 7$	7.0014	7.3594	8.1536	4.251	0.001	6
	$\bar{y}^* = 5$	4.9503	5.9922	6.3478	4.0392	0.01	7
3	$\bar{x}^* = 7$	7.0000	8.2266	8.3695	4.2747	0.001	7
	$\bar{y}^* = 5$	5.0033	8.5156	8.6391	4.3837	0.001	6

It should be noted that the search procedure only identifies design points expected to meet one of the two limits in the worst-case; checks are not performed to guarantee that the design point itself meets both limits, such as is observed for the design points for model 1 and 2 chosen to meet \bar{x}^* , which violate \bar{y}^* . (Adaptation of this approach to an effect defining the range of allowable designs meeting both limits, at least nominally, would require the search procedure to take this into account.)

The search procedure was performed on the seven identified ranges; fewer than 7 iterations were necessary to determine each limiting design point, which are listed in the fourth column of Table 3. The third column in the table indicates how close the predicted worst-case performance level is to the design limit. The expected nominal performance level, or the center of the ellipse, was defined by the information in the fifth and sixth columns. Details about the search procedure itself, including the tolerance used to identify the limiting design point and the number of iterations required to find it, are in the final two columns. It should be noted that the search tolerances were selected on a model-specific basis to produce accurate estimates without entering infinite loops.

E. Evaluating the methodology

To evaluate the use of the uncertainty characterization in approximating the limiting design points, controllers were designed for a mid-life engine at each point listed in Table 3. The approach outlined in Section II.B was applied to construct ellipses using data from Monte Carlo simulation of each closed-loop system. These were compared to the boundaries predicted by the uncertainty characterization to determine how close the worst-case performance level at each design point came to the design limit. This comparison was done graphically for models 1 to 3 through Figures 7 to 9, respectively, where the actual (solid lines) and estimated (dot-dashed lines) regions were overlaid on the simulation data (dots). Numerical comparison of the ellipses was done through the information provided in Table 4. Also indicated in the figures are the actual performance levels of the new, mid-life, and end-of-life engines at each design point (six-point stars, five-point stars, and triangles, respectively), the estimated nominal (mid-life) design points (diamonds), and the extrema of the ellipses (boxes). The design limits, $\bar{x}^* = 7$ and $\bar{y}^* = 5$, are shown as black dotted lines. These results are plotted with on the new, mid-life, and end-of-life performance level trade-offs at the trial design points (stars connected by dotted lines). Data and results for limiting design points expected to meet the limit $\bar{x}^* = 7\%$ are shown in blue, while those for points expected to meet $\bar{y}^* = 5$ seconds are shown in red (and green, for model 1).

Based on the number of points falling inside the approximated ellipses (the seventh column in Table 4), the uncertainty characterization was relatively accurate in approximating the Monte Carlo results. The most notable discrepancy between the estimated and actual boundaries was seen for the design point identified to meet the limit on \bar{x}^* for model 3, where the estimated ellipse defined a region in which under 97% of the simulation results fell (compared to 98% and above for the other test cases). A closer look at Figure 9, and the trial design point data in Figure 4, showed that this was due to the rapid change in the shape of the data distribution on the search region. As mentioned in Section III.D, a ‘peak’ formed in the distribution as the minimum HPC SM at the design point increased, leading to an increase in r_{yt} ; for model 3, this artifact became more pronounced on the search region, between 8% and 9% minimum HPC SM. Because a curve fit was used to estimate the shape of the region, the radius of the top half of the ellipse tended to be underestimated for design points closer to 8% minimum SM (such as, the limiting design point expected to meet $\bar{x}^* = 7\%$). As the design point moves closer to 9%, the estimated r_{yt} becomes more accurate, as was the case for the limiting design point expected to meet $\bar{y} = 5$ seconds.

To better evaluate results, the location of the boundaries were compared, with particular focus on how

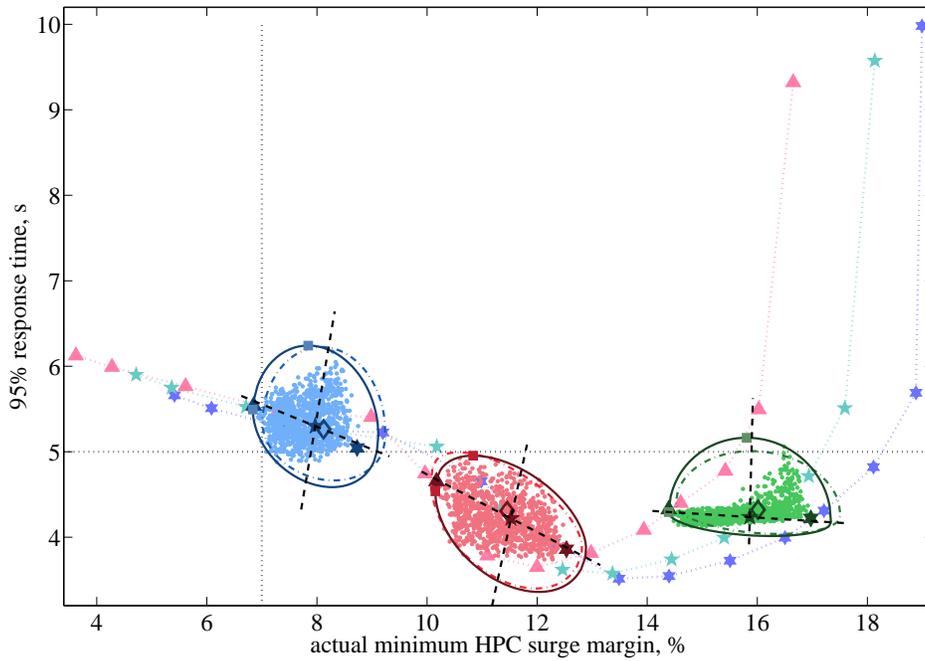


Figure 7. Plot showing how well the elliptical boundaries calculated for the test Monte Carlo simulation data match the bounds estimated using the search procedure described in Section II.D for model 1. Monte Carlo data is overlaid by the estimated (dot-dashed) and actual (solid) ellipses; axes (dashed) and extrema (boxes) for the actual boundaries are also shown. The trade-off between tr and $minSMa$ for the new (six-pointed stars), mid-life (five-pointed stars), and end-of-life (triangles) at the trial design points are indicated, along with the design limits (dotted), for reference.

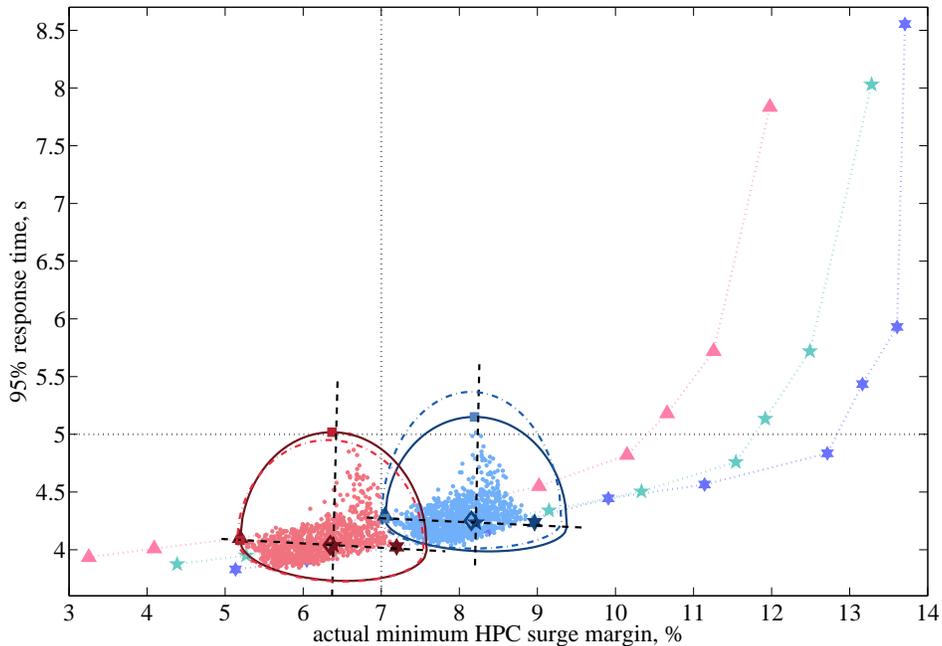


Figure 8. Plot showing how well the elliptical boundaries calculated for the test Monte Carlo simulation data match the bounds estimated using the search procedure described in Section II.D for model 2. Monte Carlo data is overlaid by the estimated (dot-dashed) and actual (solid) ellipses; axes (dashed) and extrema (boxes) for the actual boundaries are also shown. The trade-off between tr and $minSMa$ for the new (six-pointed stars), mid-life (five-pointed stars), and end-of-life (triangles) at the trial design points are indicated, along with the design limits (dotted), for reference.

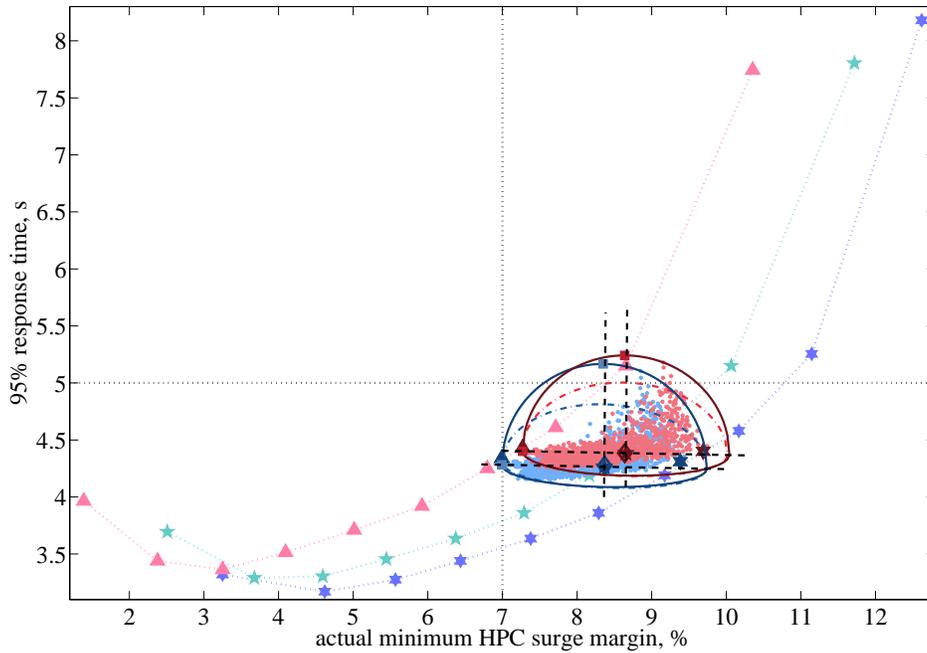


Figure 9. Plot showing how well the elliptical boundaries calculated for the test Monte Carlo simulation data match the bounds estimated using the search procedure described in Section II.D for model 3. Monte Carlo data is overlaid by the estimated (dot-dashed) and actual (solid) ellipses; axes (dashed) and extrema (boxes) for the actual boundaries are also shown. The trade-off between tr and $minSMa$ for the new (six-pointed stars), mid-life (five-pointed stars), and end-of-life (triangles) at the trial design points are indicated, along with the design limits (dotted), for reference

close the nominal and worst-case performance levels matched the estimates. This comparison is provided in Table 5, where the third and fourth columns compare the extrema of interest to the design objective (second column) and the sixth through tenth compare the estimated and actual performance level of the nominal system. Overall, the distance between the estimated and actual nominal performance levels (the last column in the table) was largest for model 1, a result that was expected since the RMS errors for $minSMa(minSMd)$ and $tr(minSMa)$ were largest for this model. It was interesting to note, however, that while the individual estimates of both $minSMa$ and tr were most inaccurate for model 1, those for the design point coordinates for model 3 were not always the most accurate, despite the distance between predicted and actual design points being smallest for this model. In particular, for the limiting design points expected to meet \bar{y}^* , model 2 predicted the response time more accurately due to the slope of the trade-off curve local to the design point. Response time at the respective points changed quicker with respect to minimum HPC SM for model 3, so that a large error in tr resulted despite an error in $minSMa$ that was three times smaller. The errors in design point location propagated to the estimates of the ellipse parameters as shown in Table 4; the large errors in r_{yt} and r_{yb} in particular were due to the shape of the surface used to estimate the boundary, as addressed in Section III.D.

For both the estimated and predicted ellipses at each limiting design point, $\psi < 0$, which suggested that errors in r_{yb} (r_{yt}) affected the limiting design point found through the search to meet the limit on \bar{x}^* (\bar{y}^*)

Table 4. Error in estimation of elliptical bounds for limiting design point listed in Table 3.

Model	Objective	r_x error (%)	r_{yt} error (%)	r_{yb} error (%)	ψ error (deg)	% covered by predicted	% covered by actual
1	$\bar{x}^* = 7$	-3.1066	3.7453	-14.5876	7.8724	99.8	99.9
	$\bar{y}^* = 5$	1.2596	-22.9415	4.3257	-1.8678	98.2	99.8
		1.5285	-35.8377	22.3886	-0.4354	99	99.2
2	$\bar{x}^* = 7$	-1.163	18.0149	-2.705	1.5527	99.9	99.9
	$\bar{y}^* = 5$	0.07519	-7.3915	1.254	-0.2578	99.6	99.8
3	$\bar{x}^* = 7$	-0.2847	-67.7329	8.4194	-0.5271	96.7	99.5
	$\bar{y}^* = 5$	-0.8293	-38.269	1.0516	-0.2063	98.2	99.5

Table 5. Actual (sixth to ninth columns) and worst-case (third and fourth columns) performance levels calculated from the Monte Carlo test data for controllers designed at the limiting points listed in Table 3; errors are calculated with respect to the predicted values for the limiting design points (Table 3). The data in the sixth and seventh column indicates the center of the ellipses drawn in the respective figure (Figures 7 to 8).

Model	Objective	Limit		$minSMd$	$minSMa$ (%)		tr (sec)		distance between centers
		Actual	Error (%)		Actual	Error (%)	Actual	Error (%)	
1	$\bar{x}^* = 7$	6.8359	2.3889	7.8047	7.9691	1.8668	5.3	-0.7183	0.1562
	$\bar{y}^* = 5$	4.9541	0.7851	10.0625	11.5241	-0.6094	4.22	2.0404	0.1122
		5.1637	-3.1214	14.7813	15.863	0.9312	4.235	2.0016	0.1724
2	$\bar{x}^* = 7$	7.0475	-0.6584	7.3594	8.2125	-0.7224	4.235	0.3764	0.06103
	$\bar{y}^* = 5$	5.0182	-1.3716	5.9922	6.3863	-0.6065	4.04	-0.01981	0.03851
3	$\bar{x}^* = 7$	6.9915	0.1214	8.2266	8.3651	0.05257	4.265	0.2269	0.01065
	$\bar{y}^* = 5$	5.2414	-4.7589	8.5156	8.6555	-0.1898	4.385	-0.02966	0.01645

(Equation (4)). The results in Table 4 suggested that the worst-case minimum HPC SM for the limiting design point meeting $\bar{x}^* = 7\%$ for model 1 was farther from this limit than in the other two models, since the errors in r_x and r_{yb} were twice as large. This was confirmed by the results in Table 5, which showed that the worst-case minimum HPC SM at this design point exceeded the limit by 2.3889% for model 1, 3.63 times more than model 2 fell short of it, and 19.68 times more than model 3 exceeded it.

Similarly, errors in r_{yt} for the design points identified to meet \bar{y}^* in models 1 and 3 were between 3 and 6 times larger than for model 2, suggesting that the worst-case limit was met most closely for model 2. With the exception of one of the design points for model 1, this expectation held: model 2 exceeded the limit $\bar{y} = 5$ seconds by 1.3716%, in the worst-case, 2.27 times less than the design point at $minSMd = 14.7813\%$ for model 1 and nearly 3.5 times less than at the respective design point for model 3. This design point proved to be least-accurate, overall, in meeting the limit it was expected to, exceeding \bar{y}^* by 4.7589%. (The design point at $minSMd = 10.0625\%$ for model 1 was most accurate, underestimating the limit by 0.7851%, likely due to the fact that the data distribution was better approximated by an ellipse than others.)[†]

F. Extension to maximum HPT temperature

This characterization of uncertainty can be extended to analyze the limiting behavior of other engine variables related to the performance and operability of the engine; these may include maximum HPT entrance temperature, minimum LPC SM, or minimum fuel-to-air ratio in the burner. Here, it was of interest to study the relationship between the maximum HPT temperature and minimum HPC SM. The steps of the algorithm described in Sections II.B and II.C were applied to the Monte Carlo data at the trial design points to determine the elliptical boundaries and to construct polynomials relating $minSMa$ to $maxT40$, and $minSMa$ and $maxT40$ to the ellipse parameters. Application of these curve fits to the design points in Table 3, provided the estimated nominal and worst-case results listed in Table 6.

Ellipses were then fit to the Monte Carlo simulation data collected at the limiting design points to evaluate the estimates; Figures 10 to 12 provide visual comparison for models 1 to 3, respectively. In these figures, the actual ellipses are shown as solid lines and the estimated ellipses as dashed lines. It is immediately clear

[†]These calculations were based on the vertical maximum of the ellipses fit to the data; the actual worst-case response time was less than this extremum for all four design points, exceeding the 5 second limit by less than 3.6% (model 3), a small improvement over the performance level estimated using r_x , r_{yt} , r_{yb} , and ψ .

Table 6. Estimated ellipse center and vertical maxima for the relationship between minimum HPC SM and maximum T_{40} at the limiting design points in Table 3.

Model	Objective	Predicted limit	$minSMa$ (%)	$maxT40$ ($^{\circ}R$)
1	$\bar{x}^* = 7$	3268.1	8.1207	3197.2877
	$\bar{y}^* = 5$	3204	11.4543	3116.1271
		3214	16.0121	3129.8097
2	$\bar{x}^* = 7$	3090.7	8.1536	3030.8418
	$\bar{y}^* = 5$	3110	8.3695	3143.8279
3	$\bar{x}^* = 7$	3241.5	6.3478	3040.4813
	$\bar{y}^* = 5$	3241.3	8.6391	3140.0752

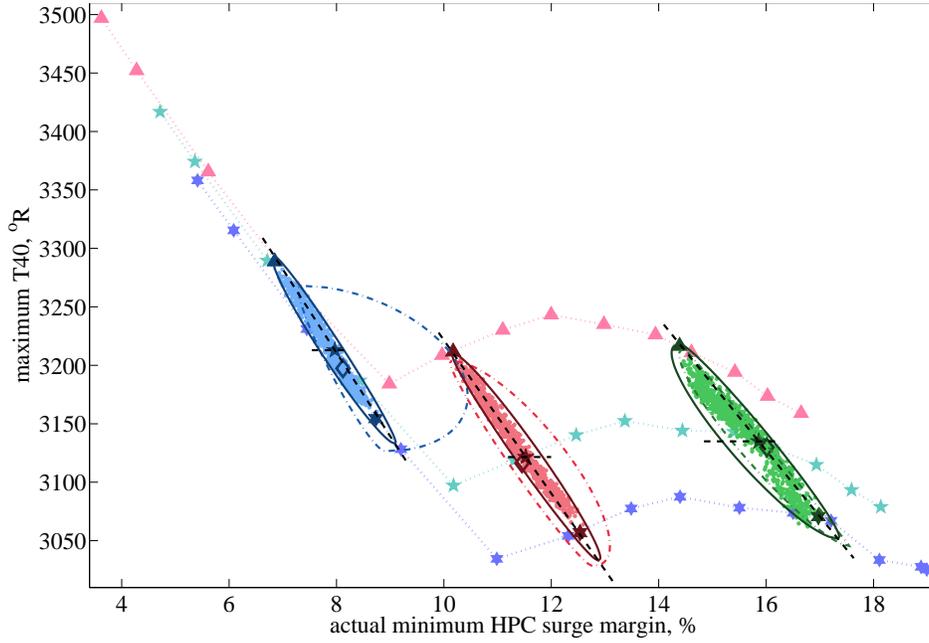


Figure 10. Plots showing the relationship between maximum T_{40} and minimum HPC SM for the three test cases for model 1 in Table 3. Also shown are the boundaries of the data estimated using curve fits (dot-dashed) and from the Monte Carlo data (solid), and the trade-off between $minSMa$ and $maxT40$ for new (six-pointed stars), mid-life (five-pointed stars), and end-of-life (triangles) at the trial design points.

that the boundaries were not as accurately estimated as when the characterization was applied to the trade-off between tr and $minSMa$. This was the case for models 1 and 3, in particular, because of inconsistent data distributions for both models. Unlike model 2, where each data distribution closely approximated an ellipse, one or more distributions for both models 1 and 3 were spread out, producing large changes in r_{yt} or r_{yb} between consecutive design points because the algorithm attempted to include as many points in the region as possible. These local extrema were captured in the curve fits and introduced errors into the approximations for design points in certain ranges. For instance, in model 1, the distribution of data collected for $minSMd = 9\%$ lay in two distinct groups, while distributions for all the other design points formed ellipses of similar size; note that the simulation data showing this relationship for the trial design points are not presented in this paper for compactness. Consequently, the ellipse fit to the data for $minSMd = 9\%$ had a much larger r_{yt} than the other trial design points and the curve fit $r_{yt}(minSMa, maxT40)$ contained a maximum to account for this. This produced estimates of r_{yt} at design points between 8% and 10% that tended to be larger than the actual value, particularly as the design point neared a minimum HPC SM of 9%. Two of the three limiting design points for this model were in or near this range, producing the inaccuracies shown in Figure 10. In model 3, similar spreading of the data occurred at $minSMd = 4\%$ and for $minSMd$ between 6% and 9%; the design points for this model were in the latter region, yielding the observed poor estimates of r_{yt} .

It was of more interest to evaluate this relationship through estimates of the vertical maximum of the ellipse, the expected maximum HPT temperature under worst-case aging. This particular extremum is important because an upper limit on HPT temperature, T_{40} , is often imposed to avoid premature structural damage from operating at too high of a temperature. The accuracy of this prediction, and the prediction of the nominal relationship, was verified by the information in Table 7. The worst-case maximum HPT temperature for models 1 and 3 was estimated within $25^{\circ}R$, an error around 0.76% (with respect to the predicted value), with much better results for models 2 (under 0.43% error). The location of the center of the ellipse was even more accurately approximated, as the actual and estimated nominal $maxT40$ differed by around $15^{\circ}R$ for model 1, and by less than $1.6^{\circ}R$ for models 2 and 3. These results suggest that most of the inaccuracy in approximating the vertical maximum for model 1 was due to the curve fit $maxT40(minSMa)$ and, for models 2 and 3, was due to the curve fits for r_x , r_{yt} and ψ (since $\psi < 0$, r_{yb} did not contribute to calculations of the maximum). The relative accuracy with which the vertical maxima at each design point were predicted using the curve fits suggests additional means for analysis and comparison of engine

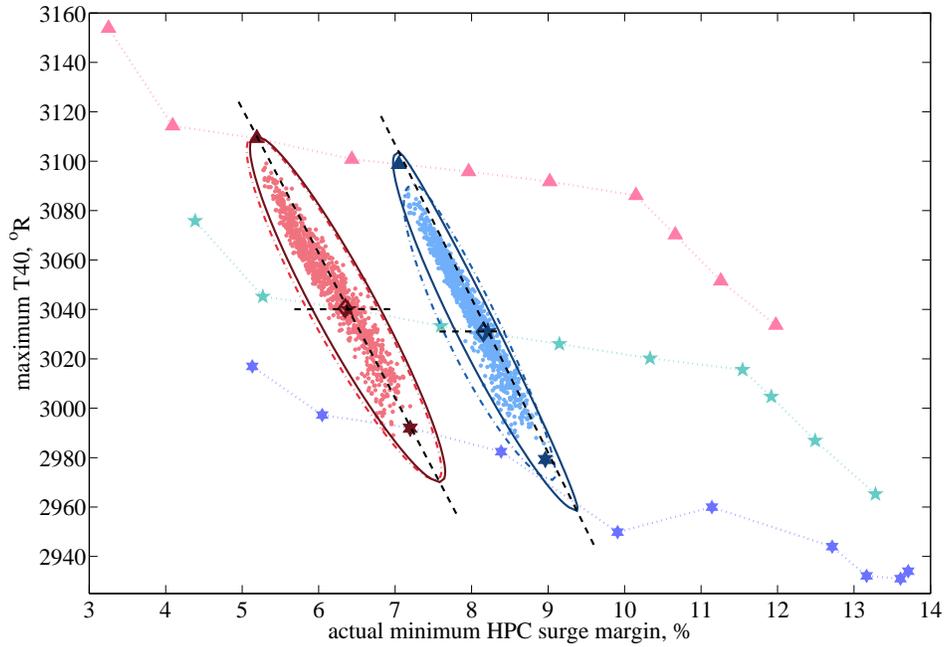


Figure 11. Plots showing the relationship between maximum T_{40} and minimum HPC SM for the three test cases for model 2 in Table 3. Also shown are the boundaries of the data estimated using curve fits (dot-dashed) and from the Monte Carlo data (solid), and the trade-off between $minSMa$ and $maxT_{40}$ for new (six-pointed stars), mid-life (five-pointed stars), and end-of-life (triangles) at the trial design points.

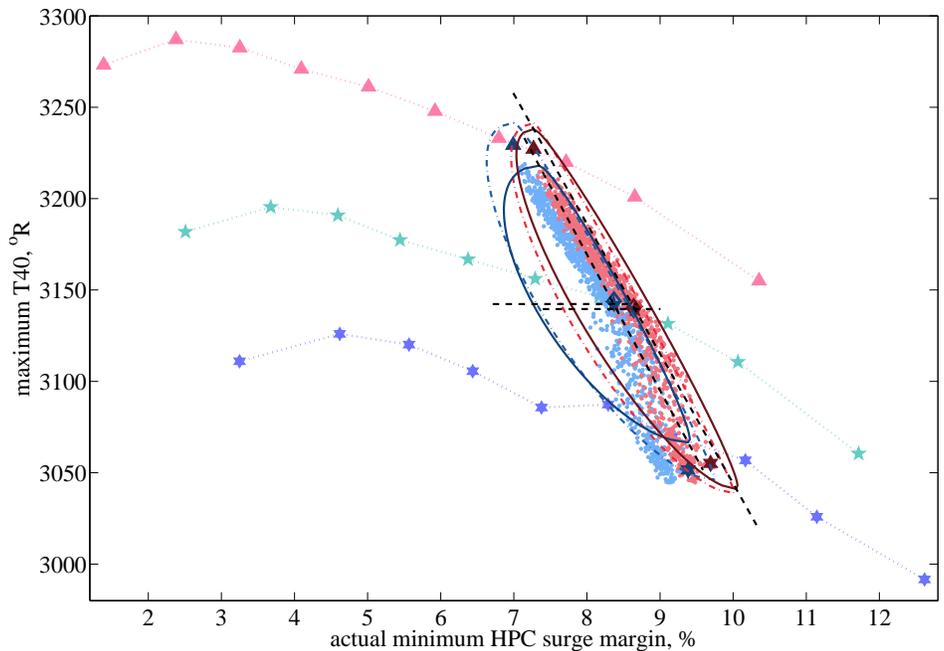


Figure 12. Plots showing the relationship between maximum T_{40} and minimum HPC SM for the three test cases for model 3 in Table 3. Also shown are the boundaries of the data estimated using curve fits (dot-dashed) and from the Monte Carlo data (solid), and the trade-off between $minSMa$ and $maxT_{40}$ for new (six-pointed stars), mid-life (five-pointed stars), and end-of-life (triangles) at the trial design points.

Table 7. Actual ellipse center (fifth to eighth columns) and vertical maxima (third and fourth columns) for the relationship between minimum HPC SM and maximum T_{40} , calculated from the Monte Carlo test data, at the limiting points listed in Table 3. Errors are calculated with respect to the estimates in Table 6.

Model	Objective	Limit		$minSMa$ (%)		$maxT40$ ($^{\circ}R$)	
		Actual	Error (%)	Actual	Error (%)	Actual	Error (%)
1	$\bar{x}^* = 7$	3292.881	-0.7583	7.9691	1.8668	3212.9038	-0.4884
	$\bar{y}^* = 5$	3210.2174	-0.1941	11.5241	-0.6094	3121.5677	-0.1746
		3218.217	-0.1312	15.863	0.9312	3134.952	-0.1643
2	$\bar{x}^* = 7$	3103.7523	-0.4223	8.2125	-0.7224	3031.0989	-0.008483
	$\bar{y}^* = 5$	3110.0769	-0.002473	6.3863	-0.6065	3040.0861	0.0130
3	$\bar{x}^* = 7$	3217.9969	0.7251	8.3651	0.05257	3142.3008	0.04857
	$\bar{y}^* = 5$	3238.0621	0.09990	8.6555	-0.1898	3139.5757	0.01591

models and may also provide another condition for consideration in selecting a limiting design point. This uncertainty characterization may be further extended to other limiting values if similar accuracy can be demonstrated for those relationships, a determination that should be done on an engine-by-engine basis.

G. Comparing the models

Three metrics were defined for comparison of the performance and operability trade-off and robustness due to aging between engine models; these are illustrated in Figure 13 for a representative pair of design points:

1. The (signed) Euclidean distance from the nominal design point to the limit for which it was chosen (for example, how far the nominal response time was from \bar{y}^* for a design point chosen to satisfy that limit in the worst-case); this is the purple dashed line in the figure
2. The (signed) Euclidean distance from the nominal design point to the limit for which it was not chosen (for example, how far the nominal minimum HPC SM was from \bar{x}^* at a design point chosen to satisfy the limit \bar{y}^* in the worst-case); this is the yellow dashed line in the figure
3. The Euclidean distance from the nominal design point to the worst-case performance level, (\bar{x}^*, \bar{y}^*) ; this is the aqua dashed line in the figure

Each metric measured something different about the engine. The first quantified how much performance level varied in the direction of the design limit; a system that was more robust to engine aging effects was expected to have a smaller value for this metric than a model that changed greatly throughout the engine life cycle. The shape of the trade-off between performance and operability was suggested, in part, by metric 2, which measured how close the nominal design point was to the limit for which it was not chosen; for an engine for which response time increases rapidly with minimum HPC SM, this metric may be expected to approach 0. The third metric provided an idea of how close the design point was to the worst-case performance level defined by \bar{x}^* and \bar{y}^* and combined information from metrics 1 and 2 into a single calculation. For an ‘ideal’ closed-loop engine model, which has no variation due to aging, and limits \bar{x}^* and \bar{y}^* chosen so that (\bar{x}^*, \bar{y}^*) is an achievable performance level, each metric would exactly equal 0.

Table 8 lists these metrics for the design points in Table 5, where the negative values of metric 2 indicated points that exceeded the limit for which the metric was calculated. Several conclusions about how the three models compare to each other were drawn from these metrics. In the ensuing discussion, models and metrics are referred to by the numeric designations in Table 8, while design points are referenced with respect to the limit they were chosen to meet in the worst-case.

The shape of the performance and operability trade-off was inferred by considering all three metrics together. Metric 3 provided the distance from the design point to the worst-case performance level, but required the information available from metrics 1 and 2 to determine where on the curve this metric held. The general trade-off trend was seen through metrics 1 and 2, which varied inversely with each other: a design point closer to one limit was generally farther from the other, as mentioned in Section II. Metrics 2 and 3, together, further suggested that the trade-off was more sensitive to changes in minimum HPC SM for models 2 and 3 and that model 1 admitted more allowable designs, since the design points meeting \bar{y}^* for this model were much farther from \bar{x}^* . This summarizes the trade-off curves in Figures 7 to 9, where the

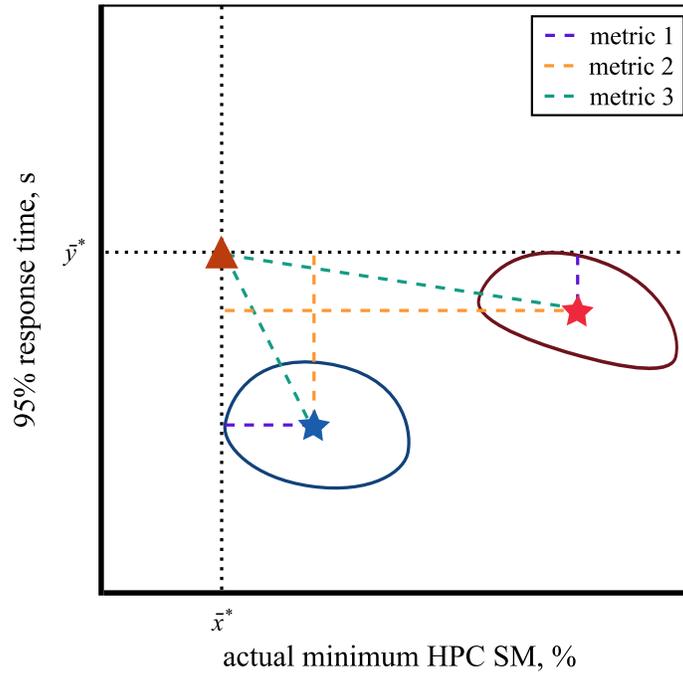


Figure 13. Illustration of metrics for a pair of generic design points (stars). Dotted lines indicate the limits \bar{x}^* and \bar{y}^* , the design point and boundary shown in blue were identified to meet \bar{x}^* and those shown in red to meet \bar{y}^* . Metric 1 is the length of the purple dashed line, metric 2 of the yellow dashed line, and metric 3 of the aqua dashed line. Note that the ‘direction’ of measurement for metrics 1 and 2 depend on the limit for which the design point was found.

presence of windup in model 1 acted to extend the range of design points meeting one of the limits beyond that for the other two models.

Because this characterization was developed to consider a single worst-case design limit, robustness of the closed-loop model was best represented through metric 1, which measured how much minimum HPC SM or response time varied throughout the life of the engine. This provided a way to summarize the shape of the ellipse at the design point. A model more robust to aging in the given objective *at the specified design point*, such as model 1 for the limit \bar{x}^* and model 3 for the limit \bar{y}^* , was expected to have a smaller value for metric 1. This was confirmed through the results in Figures 7 to 9, where the Monte Carlo simulation results for the design points meeting $\bar{x}^* = 7\%$ were less spread out along the minimum HPC SM direction (x -axis) for model 1, and those for points meeting $\bar{y}^* = 5$ seconds were less spread out along the direction of response time (y -axis) for model 3. The metrics did not provide information about robustness at other design points, since data at those points were not part of the calculations.

H. Summary of results

When designing a controller for an engine model, it is often of interest to choose specifications that guarantee a desired level of performance in the presence of worst-case uncertainty. Considering the effect of aging on a closed-loop engine, an approach for estimating the design points at which a specific performance level limit is met was developed. This involved characterizing the uncertainty through an ellipse that defined the worst-case performance level expected at a particular design point, using Monte Carlo simulation results.

Table 8. Metrics calculated for the three example models, using the design points determined through Monte Carlo simulation instead of using the curve fit estimations. Recall that model 1 has two design points meeting \bar{y}^* , while the other models only have a single point meeting this limit, as shown in Figures 7 to 8 and indicated in Tables 3 and 5.

Objective	Metric 1			Metric 2			Metric 3		
	Model 1	Model 2	Model 3	Model 1	Model 2	Model 3	Model 1	Model 2	Model 3
$\bar{x} = 7$	0.9691	1.2125	1.3651	-0.3	0.765	0.735	1.0145	1.4337	1.5504
$\bar{y} = 5$	0.78	0.96	0.615	4.5241	-0.6137	1.6555	4.5908	1.1394	1.7660
	0.765			8.863			8.896		

A series of curve fits was constructed from the parameters of ellipses calculated for several trial design points, and was incorporated in a binary search procedure to identify those design points meeting a prescribed worst-case limit on minimum HPC SM or response time.

The method was evaluated using three variations of the C-MAPSS40k engine model in closed-loop with a controller designed by TTECTrA. Each engine model was defined by the same maps, but had a different set of scalars to size the engine and affect its response. Controllers were designed at each point identified by the search and data were collected from Monte Carlo simulations of each closed-loop system. The worst-case performance level at each design point was found to be within 4.76% of the design limit. Inaccuracy in the curve fits, particularly for model 1, and inconsistencies in the shape of the data distributions at the design points likely produced most of the discrepancy. The approach was extended to the relationship between minimum HPC SM and maximum HPT temperature to evaluate whether estimation of the worst-case maximum HPT temperature was possible. Although the shape of the elliptical regions in which data lay on the plane was poorly approximated, the vertical maxima were estimated within 0.76% ($25^\circ R$), suggesting that additional limiting values may be estimated by the characterization to provide another point of analysis for the closed-loop system or an additional constraint in determining the limiting design point.

The three models were compared through metrics that provided information about the shape of the trade-off between minimum HPC SM and response time and the robustness of the closed-loop system due to aging. Most notably, these metrics provided a summary of the Monte Carlo simulation results, which showed that model 1 provided a larger range of allowable design points than models 2 and 3, due to response time being more affected by windup and less sensitive to increases in minimum HPC SM. At the design points identified for models 2 and 3, the metrics also indicated increased variation in response time due to aging in comparison to model 1, which was also observed in the simulation data.

Two design points did not fall within the range of acceptable performance levels due to the simplicity of the search procedure implemented to identify these points. Modification to the search may include ensuring the limiting design points nominally meet both limits on \bar{x}^* and \bar{y}^* , or that both limits are met in the worst-case. The latter would enable identification of models that admit few or no design points. The search may also be expanded to consider other limits related to performance and operability of the engine, such as maximum HPT temperature. At a higher level, implementation of this uncertainty characterization may be streamlined through more efficient data collection to allow it to be integrated in a future release of TTECTrA, expanding the capability of the tool by automating the identification of limiting design points.

IV. Conclusion

The proposed characterization of uncertainty may be used to study and compare the trade-off between response time and minimum high-pressure compressor (HPC) surge margin (SM) in engine models. When designing a controller to meet either a minimum HPC SM, \bar{x}^* , or maximum response time, \bar{y}^* , throughout the engine life, a binary search utilized the results of the characterization to identify design points at which worst-case performance level exceeded the limit by less than 4.76%. Additionally, extension of the characterization to other engine parameters related to performance and operability was demonstrated for the maximum high-pressure turbine (HPT) entrance temperature, $maxT40$; it was shown to provide accurate estimation of the worst-case $maxT40$ (within 0.76%), a result that may be used in analysis of the system or to refine the design point identification process. Such refinement may also include implementing the search procedure in such a way to guarantee that both \bar{x}^* and \bar{y}^* are met nominally and throughout the engine life-cycle. The metrics that were defined reduced the information available from plots into three numerical values. Together, these metrics represent the shape of the performance and operability trade-off and the robustness of the system (at the design point) and may be used to compare different engines. For the models considered here, the metrics suggested that the most robust systems are model 1, for design meeting a minimum HPC SM limit, and model 3, for design meeting a maximum response time limit, and that the performance level of model 1 is less sensitive to the minimum HPC SM than the other models. Depending on design requirements, this may help guide the engine design itself by providing information on the feasible control designs, and the resulting dynamic performance, possible for the engine. Additionally, the results provide improved confidence that the transient requirements for a given design are met throughout the life cycle of the engine.

References

- ¹Csank, J. T. and Zinnecker, A. M., "Application of the Tool for Turbine Engine Closed-Loop Transient Analysis (TTECTrA) for Dynamic Systems Analysis," *AIAA Propulsion and Energy Forum and Exposition 2014: 50th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, AIAA-2014-3975, Cleveland, OH, July 2014.
- ²Green, M. and Limebeer, D. J. N., *Linear Robust Control*, Dover Publications, Inc., Mineola, NY, 1995.
- ³Zhou, K., Doyle, J. C., and Glover, K., *Robust and Optimal Control*, Prentice Hall, Inc., Upper Saddle River, NJ, 1996.
- ⁴Zhou, K. and Doyle, J., *Essentials of Robust Control*, Prentice Hall, Inc., Upper Saddle River, NJ, 1998.
- ⁵Garg, S., "Robust Integrated Flight/Propulsion Control Design of a STOVL Aircraft Using H-infinity Control Design Techniques," *Automatica*, Vol. 29, No. 1, 1993, pp. 129–145.
- ⁶Frederick, D., Garg, S., and Adibhatla, S., "Turbofan engine control design using robust multivariable control technologies," *Proceedings of the 32nd AIAA, ASME, SAE, and ASEE Joint Propulsion Conference and Exhibit*, AIAA 96-2587, Lake Buena Vista, FL, July 1996.
- ⁷Kantner, M. J., *Robust Simulation and Analysis of Nonlinear Systems*, Ph.d. thesis, California Institute of Technology, May 1997.
- ⁸Lu, W.-M. and Doyle, J. C., "Robust Analysis and Synthesis for Nonlinear Uncertain Systems," *IEEE Transactions on Automatic Control*, Vol. 42, No. 12, December 1997, pp. 1654–1662.
- ⁹Fekri, S., Bates, D. G., and Postlethwaite, I., "Linear vs. Nonlinear Robustness Analysis: A Case Study," *Proceedings of the 16th IEEE International Conference on Control Applications (Part of the IEEE Multi-conference on Systems and Control)*, TuB04.2, Singapore, October 2007, pp. 753–758.
- ¹⁰Kenny, S. P., Crespo, L. G., Andrews, L., and Giesy, D., "Robust Control Design for Uncertain Nonlinear Dynamic Systems," *Proceedings of IMAC 30 - A Conference and Exposition on Structural Dynamics*, Jacksonville, FL, January-February 2012.
- ¹¹May, R. D., Csank, J., Lavelle, T. M., Litt, J. S., and Guo, T.-H., "A High-Fidelity Simulation of a Generic Commercial Aircraft Engine and Controller," *Proceedings of the 46th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, AIAA-2010-6630, Nashville, TN, July 2010.
- ¹²Csank, J., May, R. D., Litt, J. S., and Guo, T.-H., "Control Design for a Generic Commercial Aircraft Engine," *Proceedings of the 46th AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, AIAA-2010-6629, Nashville, TN, July 2010.
- ¹³Csank, J. T. and Zinnecker, A. M., *Tool for Turbine Engine Closed-Loop Transient Analysis (TTECTrA) User's Guide*, NASA TM 2014-216663, June 2014.
- ¹⁴"Part 14. Aeronautics and Space," Federal Aviation Administration, Electronic Code of Federal Regulations, http://www.ecfr.gov/cgi-bin/text-idx?SID=6ce275c613bbaa9200e6c9c968878bdc&node=se14.1.33_173&rgn=div8 [last accessed January 14, 2015].