AIAA SciTech Forum
9 - 13 January 2017, Grapevine, Texas
58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference

Genetic Algorithm-Guided, Adaptive Model Order Reduction of Flexible Aircrafts

Jin Zhu[1], Yi Wang[2], Kapil Pant[3]
CFD Research Corporation, Huntsville, AL 35806

and

Peter Suh[4], Martin J. Brenner[5]
NASA Armstrong Flight Research Center, Edwards, CA 93523

This paper presents a methodology for automated model order reduction (MOR) of flexible aircrafts to construct linear parameter-varying (LPV) reduced order models (ROM) for aeroservoelasticity (ASE) analysis and control synthesis in broad flight parameter space. The novelty includes utilization of genetic algorithms (GAs) to automatically determine the states for reduction while minimizing the trial-and-error process and heuristics requirement to perform MOR; balanced truncation for unstable systems to achieve locally optimal realization of the full model; congruence transformation for "weak" fulfillment of state consistency across the entire flight parameter space; and ROM interpolation based on adaptive grid refinement to generate a globally functional LPV ASE ROM. The methodology is applied to the X-56A MUTT model currently being tested at NASA/AFRC for flutter suppression and gust load alleviation. Our studies indicate that X-56A ROM with less than one-seventh the number of states relative to the original model is able to accurately predict system response among all input-output channels for pitch, roll, and ASE control at various flight conditions. The GA-guided approach exceeds manual and empirical state selection in terms of efficiency and accuracy. The adaptive refinement allows selective addition of the grid points in the parameter space where flight dynamics varies dramatically to enhance interpolation accuracy without over-burdening controller synthesis and onboard memory efforts downstream. The present MOR framework can be used by control engineers for robust ASE controller synthesis and novel vehicle design.

Nomenclature

A	=	state matrix
B	=	input matrix
C	=	output state matrix
D	=	input transition
f_i	=	individual fitness value
F_i	=	individual objective function value
J	=	objective function
M	=	matrices in state space model
pb	=	pitch rate
$\tilde{P}$	=	generalized controllability gramian
$\tilde{Q}$	=	generalized observability gramian
qb	=	pitch rate
R	=	common subspace for reprojection
S	=	Singular value matrix
T	=	transformation matrix for consistent state representation

[1] Associate Engineer, CFD Research Corporation, non AIAA Member
[2] Director, CFD Research Corporation, AIAA Member; yi.wang@cfdrc.com
[3] Vice President, CFD Research Corporation, non AIAA Member
[4] Aerospace Engineer, Aerostructures Branch, and AIAA Member
[5] Aerospace Engineer, Aerostructures Branch, and AIAA Senior Member

u = input signals
y = response measurements
$\tilde{V}$ = transformation matrix in balanced realization
$\tilde{W}$ = transformation matrix in balanced realization
W = objective function weight
x = system state
ρ = a vector of measurable parameters
Φ = Right unitary matrix in singular value decomposition

I. Introduction

With the fast paced technological advances in this new era of science, modern aerospace designs are able to incorporate new flexible structures and lighter materials to achieve better maneuverability, endurance, and performance. As a result they are also more susceptible to issues such as complex dynamics and interactions between the controller and the aerodynamic and structural systems, which may lead to catastrophic events such as flutter, limit cycle oscillation, and gust loading. In order to design a modern flexible aircraft that can attain a safe and acceptable flight envelope, detailed modeling and high fidelity simulations of aeroservoelastic (ASE) systems must be performed prior to flight tests to be able to prevent aeroelastic (AE) failures. While full-order models coupling the nonlinear aerodynamics with structural models are capable of accurate prediction of underlying AE phenomena and onset, their prohibitive computational cost, low speed, nonlinear nature, as well as difficulty to deploy controllers with high-state-order models render it impractical for integration in the design environment involving concurrent ASE analysis and control synthesis and design.

To combat these challenges various model order reduction (MOR) techniques have been developed in the context of linear parameter varying (LPV) formulation. In LPV, the fully coupled nonlinear aircraft model is represented as an ensemble of linearized models at the grid points within the parameter space and the model parameters vary across the flight envelope. Models at any location within the domain can be obtained by interpolating those at the grid points. MOR aims to reduce the full-order LPV ASE model into a reduced state-space form while retaining the dominant dynamics of the system in the target frequency range where AE may be involved. MOR can be classified into non-transformation (e.g., truncation and residualization) and transformation-based techniques, such as modal reduction, balanced truncation, Krylov based projection, hybrid Singular Value Decomposition (SVD)–Krylov approaches, etc. While yielding reduced order models (ROMs) with consistent states amenable to direct model interpolation, the former typically is a trial-and-error and/or empirical process that manually examines and selects unimportant states to eliminate from the original systems in an iterative manner. On the other hand, for the transformation-based MOR techniques, although more efficient and accurate given tight state budgets, the optimal transformation is flight parameter-dependent leading to different state meanings of the ROMs at various parameter locations in the flight envelope, which destroys the state consistency of the LPV system and makes immediate interpolation impossible. In order to form global LPV ROMs encompassing the entire parameter space, several techniques have been proposed for consistent state representation and ROM interpolation. Hjartarson et al. [1] proposed to apply the transformation matrix obtained by balanced truncation at a single flight condition to the LPV model sets within the entire flight envelope. Although maintaining state consistency, the approach is sub-optimal as balancing transformation by nature changes with flight parameter. Moreno et al. [2] used a contractive right coprime factorization approach to attain consistent controllability and observability gramians throughout the flight envelope which can be balanced to achieve state consistency. This approach suffers from several inherent limitations, such as difficulty to use in broad and high-dimensional parameter space, and numerical issues associated with high state orders. Panzer et al. [3] proposed two methods, respectively, based on reprojection into a common subspace and optimization-based matrix matching to achieve identical state meanings among local models for interpolation. The former was employed for interpolating LPV ROMs of industrial flexible aircrafts [4]. The common subspace in [3] is obtained by the SVD of the ensemble of the transformation matrices at various parameters, and hence, is ill-suited for uses in broad flight envelope including dramatically varying parameters. Poussot-Vassal and Demourant [5] compared several MOR techniques and their applications to aircraft systems, including balanced truncation, iterative tangential interpolation algorithm, and iterative SVD-tangential interpolation algorithms along with a modal form-based coordinate transformation to achieve state consistency. Theis et al [6] developed a new modal matching technique, which casts ROMs into a mode-wise canonical form and matches modes with similar dynamic properties at neighboring grid points to minimize state inconsistency.

This paper presents the development of LPV ASE reduced order models (ROMs) of flexible aircrafts based on a combination of sequential model order reduction (MOR), consistent state representation, and model interpolation

approaches. The X-56A MUTT vehicle with flexible wings currently being tested at NASA/AFRC for flutter suppression and gust load alleviation was used for analysis, verification, and demonstration. In the sequential MOR the truncation and residualization methods were first applied to the states of sensors, actuators, aerodynamic lags, rigid bodies, elastic structures of the full-order X-56A MUTT ASE model on the grid points. In contrast to previous efforts [1, 2, 7], for the first time a genetic algorithm (GA) was used for automated selection of unimportant states in the aerodynamic lag and the elastic modes to truncate/residualize, which not only eliminates the manual and trial-and-error process for state screening, but also improves MOR accuracy and performance given state budgets. The balanced truncation for unstable systems was also utilized to further remove states with minor contribution to the input/output energy of the system in the local reduced model. Next, the method of congruence transformation was employed to remedy the issue of inconsistent state representation among the local ROMs caused by the flight parameters-dependent transformation as discussed above. Different from the previous approach of reprojection onto a common subspace [3, 4, 7], the congruence transformation allows "weak" fulfillment of the "Modal Assurance Criterion" (MAC) and ROM interpolation to construct a unified LPV ROM applicable across broader flight envelope. Finally, an adaptive grid refinement strategy was developed to selectively add grid points at the regions where system response is susceptible to the flight parameters and ensure smooth interpolation and transition of ROMs within the domain.

II. X-56A MUTT Model

The linear time invariant (LTI) state-space models of the X-56A MUTT airframe were provided by NASA/AFRC. They were developed using the generalized mass, stiffness, and aerodynamic matrices obtained by MSC/Nastran [8] and ZAERO [9]. There are 10 control surfaces on the vehicle, five on each wing; and 2 throttle controls for engine dynamics as shown in Figure 1. The five actuator inputs for control surfaces on the left wing are labeled as BFL, WF1L, WF2L, WF3L, and WF4L starting from the inner body to the outer wing tip with units of degrees. Likewise, the actuators on the right wing are labeled as BFR, WF1R, WF2R, WF3R, and WF4R based on the same convention. The rigid-body state sensors (IMU-MIDG) are located around the center of the vehicle, while the accelerometer locations are, respectively, placed at the front of the vehicle (ASESNSR100 with a channel name cfz_200_g), at the rear (ASESNSR1000), at the leading and trailing edge of the left wing (ASESNSR400 and ASESNSR600), and of the right wing (ASESNSR1100 and ASESNSR1300).

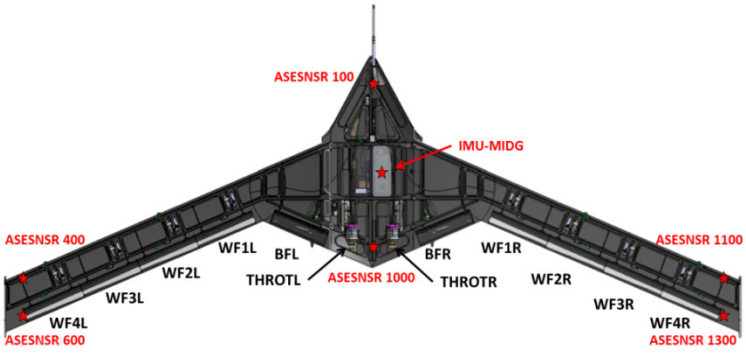

Figure 1. Sensors and actuators deployment in the X-56A MUTT vehicle

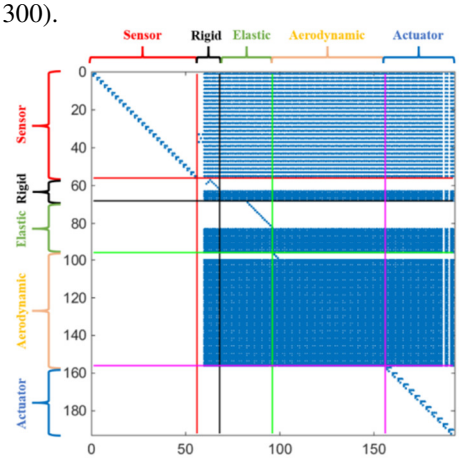

Figure 2. Sparsity pattern and partition of A matrix

A set of 99 models were generated at $M = 0.16$ on grid points of a 2D parameter space across the flight envelope. The two parameters are KEAS (knots equivalent airspeed), which ranges from 50 KEAS to 150 KEAS in 10 KEAS increments and fuel weight, which ranges from 0 lb to 80 lb in 10 lb increments. The models have 56 states corresponding to the 2nd-order sensors (28 in total), 12 rigid body states, 14 elastic structural modes and 14 derivatives (modal velocity), 60 aerodynamic lag states, and 36 states for the third order actuators (12 control surfaces). According to the V-g and V-f plots of the X-56A baseline model at $M = 0.16$ [10], the normalized flutter frequencies for SBFF (symmetric body freedom flutter), SWBTF (symmetric wing bending torsion flutter), and AWBTF (anti-symmetric wing bending torsion flutter) modes are, respectively, at 1, 3.68, and 3.912 (all the flutter frequencies are normalized by the one for SBFF). The target normalized frequency range ω for X-56A model reduction is determined to be $0.01 < \omega < 5.37$ to ensure full coverage of the flutter behavior of interest and system

response. The sparsity pattern of A matrix is illustrated in Figure 2. The physical meaning of the states and their corresponding entries in A is utilized to guide the MOR process for constructing ROMs.

Instead of using raw inputs and outputs, composite inputs and outputs constructed by the raw ones were included in the model to facilitate controller synthesis. Table 1 and Table 2 illustrate the raw and composite inputs/outputs and their associated channels as well. Specifically composite inputs 1-4 and 5-6 are, respectively, used as control means for stabilization and damping augmentation along the pitch and roll axes, and composite outputs 1 and 2-4 for the rigid-body and ASE observation along the roll and pitch axes.

Table 1. Raw and Composite inputs and related channels for actuators in X-56A MUTT ASE Model

Raw Inputs	Raw Channels	Composite Inputs	Composite Channels
WF1L & WF1R	*wf1l_cmd_deg & wf1r_cmd_deg)*	(WF1L+WF1R)/2	*(wf1l_cmd_deg + wf1r_cmd_deg)/2*
WF2L & WF2R	*wf2l_cmd_deg & wf2r_cmd_deg*	(WF2L+WF2R)/2	*(wf2l_cmd_deg + wf2r_cmd_deg)/2*
WF3L & WF3R	*wf3l_cmd_deg & wf3r_cmd_deg*	(WF3L+WF3R)/2	*(wf3l_cmd_deg + wf3r_cmd_deg)/2*
WF4L & WF4R	*wf4l_cmd_deg & wf4r_cmd_deg*	(WF4L+WF4R)/2	*(wf4l_cmd_deg + wf4r_cmd_deg)/2*
		(WF3L-WF3R+WF2L-WF2R)/4	*(wf3l_cmd_deg - wf3r_cmd_deg + wf2l_cmd_deg - wf2r_cmd_deg)/4*
		(WF4L-WF4R)/2	*(wf4l_cmd_deg - wf4r_cmd_deg)/2*

Table 2. Raw and composite outputs and related channels for sensors in X-56A MUTT ASE Model

Raw Inputs	Raw Channels	Composite Inputs	Composite Channels
p	*pb_gyro_200_dps*	*pb*	*pb_gyro_200_dps*
q	*qb_gyro_200_dps*	*qb*	*qb_gyro_200_dps*
ASESNSR400	*loft_200_g*	ASE1	*(1-.3433)*(lofz_200_g + rofz_200_g)/2 + (1-.4362)*(roaz_200_g + loaz_200_g)/2+(.3433+.4362...)*
ASESNSR1100	*rofz_200_g*		
ASESNSR1300	roaz_200_g		
ASESNSR600	loaz_200_g		
ASESNSR1000	caz_200_g	ASE2	*caz_200_g*(.72)+cfz_200_g*(.28)*
ASESNSR100	cfz_200_g		

III. Linear Parameter-Varying Aeroservoelastic Models of Aircraft

Linear parameter-varying (LPV) models are state-space models whose mathematical descriptions are functions of time-varying parameters, i.e.,

$$\begin{bmatrix} \dot{x} \\ y \end{bmatrix} = \begin{bmatrix} A(\rho) & B(\rho) \\ C(\rho) & D(\rho) \end{bmatrix} \begin{bmatrix} x(t) \\ u(t) \end{bmatrix} \tag{1}$$

where $A(\rho)$ is the state matrix, $B(\rho)$ is the input matrix, $C(\rho)$ is the output state matrix, $D(\rho)$ is the input transition matrix, $\rho \in \Re^{np}$ is a vector of measurable parameters, and specifically, is the fuel weight and KEAS in the present work, $u \in \Re^{nu}$ and $y \in \Re^{ny}$ are, respectively the vector of the control inputs and measurement outputs. There are several methods to represent the parameter dependence in LPV models above, such as linear fractional transformation, polytopic dependence of the state matrix on the parameters, linearization on a gridded domain, etc. This paper targets the MOR of LPV models based on the gridded domain to agree with the full-order X-56A models. The gridded domain LPV is illustrated in Figure 3, in which the nonlinear dynamics in the ASE system of the aircraft is treated as its linearization around various flight operating points (also termed grid points or parameter locations hereafter). A set of original, full-order Linear Time Invariant (LTI) state space models are first constructed at the grid points in the flight parameter space, and then can be used for ROM generation and controller synthesis.

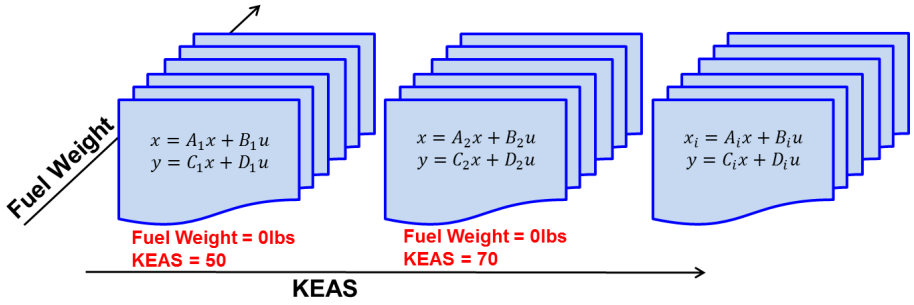

Figure 3. linear parameter varying (LPV) formulation of the aeroservoelastic (ASE) models of aircraft

IV. Model Order Reduction (MOR) for LPV ASE Models of Flexible Aircraft

Figure 4 illustrates our MOR methodology for LPV ASE models of aircraft. A prerequisite of the approach to constructing LPV ROMs is to first have a set of full-order LTI state space models describing coupled ASE and flight control behavior at grid points in the parameter space. The full model can be generated from various relevant modeling tools as shown in the blue box in Figure 4. The entire MOR process includes two steps: (1) **Local MOR**: the full-order LTI model set is first reduced and transformed onto a low-dimension subspace to generate a set of local ROMs. Several techniques can be used, including truncation and residualization, transformation and truncation (e.g., modal reduction and balanced realization and truncation, Krylov methods, and their combinations); and (2) **Model Interpolation and PV ROM Realization** (green box): the global LPV ROM applicable to the entire flight envelope is obtained by interpolating the system matrices of the local ROM set obtained in the previous step. As the transformation used in step (1) depends on the location of the grid points, measures need to be taken to ensure all the ROMs are cast in a consistent state representation (or coordinates) prior to model interpolation. Eq. (2) summarizes the MOR process

$$\begin{bmatrix}\dot{x}\\ y\end{bmatrix}=\begin{bmatrix}A_i & B_i\\ C_i & D_i\end{bmatrix}\begin{bmatrix}x(t)\\ u(t)\end{bmatrix}\xrightarrow[\text{Consistent State Represenation}]{\text{MOR}}\begin{bmatrix}\dot{x}_i^*\\ y\end{bmatrix}=\begin{bmatrix}A_i^* & B_i^*\\ C_i^* & D_i\end{bmatrix}\begin{bmatrix}x_i^*\\ u(t)\end{bmatrix}\xrightarrow{\text{LPV ROM}}\begin{bmatrix}\dot{x}^*\\ y\end{bmatrix}=\begin{bmatrix}A_{PV}^*(\rho) & B_{PV}^*(\rho)\\ C_{PV}^*(\rho) & D_{pV}(\rho)\end{bmatrix}\begin{bmatrix}x^*\\ u(t)\end{bmatrix} \quad (2)$$

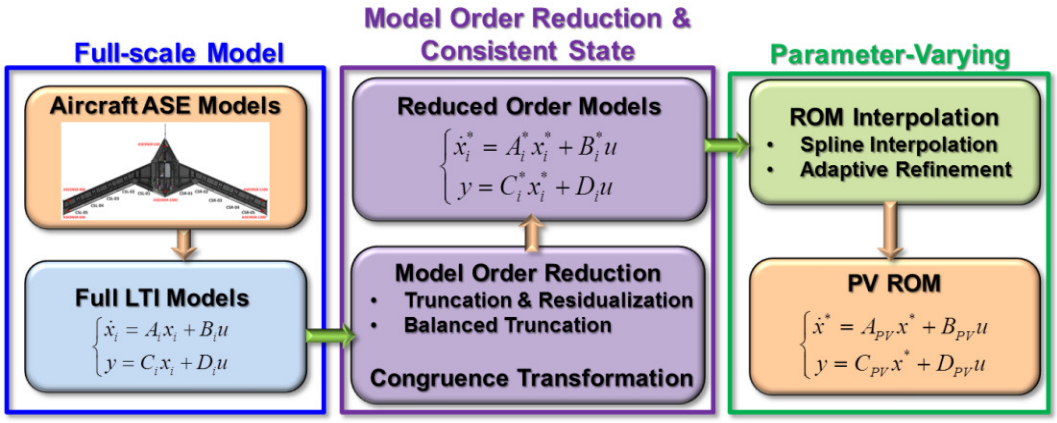

Figure 4. Organization of linear parameter varying (LPV) model order reduction (MOR) framework

Two techniques serve as the main workhorse in our framework for the local MOR process, namely, truncation and residualization and balanced truncation for unstable systems as described in the sub-sections A and C below, respectively. The truncation and residualization will be first used to remove common states of unimportance among models at all grid points without altering state consistency. It is then followed by the balanced truncation to further refine the model from the input/output channel energy perspective. There are two points of particular note: (1) rather than relying on the user's experience and trial-and-error iteration, a genetic algorithm-based procedure was developed to "intelligently" determine which states to retain (or remove) in the aforementioned truncation and residualization (see Section B below); (2) in contrast to the previous research [6, 7] including ours, the modal reduction approach based on the real and ordered eigenstructure decomposition was not used in the present effort. This is because the modal frequency varies significantly across the broad 2D flight parameter space and the full

models at various grid points have the different number of complex and real poles, and the use of modal reduction causes substantial state inconsistence among locally reduced ROMs.

A Truncation and Residualization

MOR by truncation and residualization essentially partitions the state vector x in the model into two components $[x_1\ x_2]^T$, where x_1 are the states to keep and x_2 are those to eliminate. Therefore the system matrices, A, B, and C can be partitioned as:

$$A=\begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix},\quad B=\begin{bmatrix} B_1 \\ B_2 \end{bmatrix},\quad C=\begin{bmatrix} C_1 & C_2 \end{bmatrix} \tag{3}$$

The ROM is obtained by truncating all the terms associated with x_2. Truncation preserves the ROM accuracy at high frequencies. When the steady state gain of a system needs to be retained, a residualization procedure is implemented, in which the state derivatives for x_2 are set to zero, leading to a more accurate approximation of the original system at low frequency. The residualized ROM is then given by,

$$\begin{bmatrix} \dot{x}_1 \\ y \end{bmatrix}=\begin{bmatrix} \left(A_{11}-A_{12}A_{22}^{-1}A_{21}\right) & \left(B_1-A_{12}A_{22}^{-1}B_2\right) \\ \left(C_1-C_2A_{22}^{-1}A_{21}\right) & \left(D-C_2A_{22}^{-1}B_2\right) \end{bmatrix}\begin{bmatrix} x_1(t) \\ u(t) \end{bmatrix} \tag{4}$$

B Genetic Algorithm (GA)-Guided Truncation and Residualization

The truncation and residualization above will be employed to reduce states under various categories (e.g., sensor, actuators, rigid-body, aerodynamic lag, elastic states, etc.) of the aircraft model. The challenge is to determine which states to keep (or remove). In previous efforts this is performed based on the understanding of the vehicles dynamics [2, 11], iterative process [1], retention of the states corresponding to leading modes [7], etc., which are mostly empirical. In the present effort, we developed a global optimization approach using genetic algorithm to automatically determine which states to reduce with minimal reliance on user's experience and trial-and-error process while maintaining dynamics of the original system.

Genetic Algorithm (GA) is a stochastic global search method designed to mimic evolution and natural selection. GAs uses a multitude of initial potential solutions (often called a population) to sample the entire solution space and then use the process of natural selection to evaluate the population based on the individual fitness levels of its chromosomes (member of the population) to better approximate the optimal solution [12]. This process takes many generations to converge and at each generation, a new population is created based on the fitness level of the previous generation. As shown in Figure 5, most GA applications have four key aspects: mutation, crossover, objective function, and fitness function/selection. The first step in initializing any type of GA application involves creating an initial population. A population is a set of binary strings (often called genotypes or chromosomes) each of which is a potential solution to the optimization problem. The chromosomes must be altered in some way to approach the optimal solution, which in our case is the most important states to retain. In all stochastically based algorithms, there must be a tradeoff between exploration (sampling the solution space), and exploitation (fixing the solution on the global minimum). In a GA, the role of exploration falls on the mutation operator and the role of the exploitation on the crossover operator [12]. For binary encoding, the mutation operator typically uses bit inversion to accomplish its task. Each gene is altered independently of others, typically with a probability of $1/L$, where L is bounded by the length of the population size and the length of the chromosome. The probability of mutation is usually very low as its main purpose is to maintain the diversity within the population and make sure that the solution does not prematurely converge. Mutations within a GA is similar to performing a random walk process [13, 14]. The crossover operator is used as an exploitation tool and combines portions of two different chromosomes that are already "good," with the goal to generate the offspring approaching closer to the global minimum. Only using crossover within a GA is prone to be trapped with a local minimum [13, 14]. We also devised a way to decode the binary string carried by chromosomes into a physical solution. An index of "1" indicates that the state would be kept and "0" be removed. For example, consider a system with state vector $x = [x_1\ x_2\ x_3\ x_4\ x_5]^T$ that has the binary representation [1 1 0 1 0]; this means that the reduced system has the configuration $x_r = [x_1\ x_2\ x_4]^T$ with x_3 and x_5 eliminated.

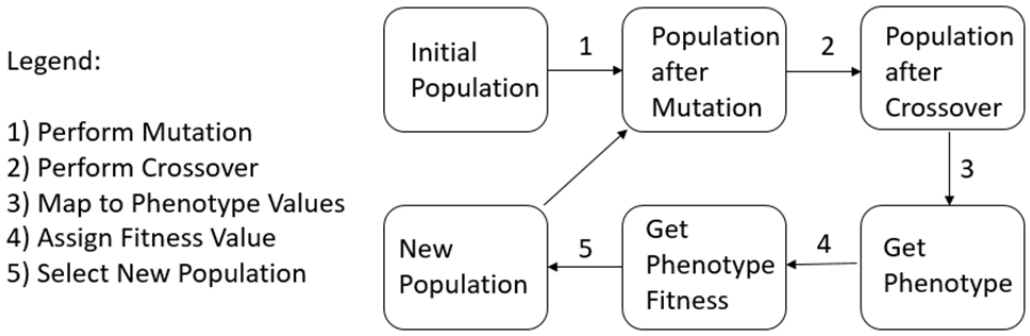

Figure 5. Flow diagram of a generic Genetic Algorithm

Similar to any optimization problem, one key aspect of GA is the objective/cost function. The goal of MOR is to reduce the number of states in the ASE system while preserving system responses along target input-output channels. A good candidate for the objective function in our case is the discrepancy in the system response between the original and the reduced system. In this work the infinity norm along with a weighting function (W) in the frequency range of interest is used

$$J = \left\| W(j\omega)\left(G(j\omega) - G_r(j\omega)\right)\right\|_\infty \tag{5}$$

where G and G_r, respectively, represent the dynamic system before and after the truncation and residualization. To evaluate the fitness level of the candidate population, a fitness functions was also defined in Eq. (6), which yields high fitness values f_i while the objective function value F_i is low

$$f_i = \left(c/F_i\right)^n \tag{6}$$

where c and n are constant weights. Once the fitness value is obtained, the roulette wheel method is used to select which genomes move onto the next generation. In this method, each genome is assigned a probability of survival based on its individual fitness value, viz.,

$$P_i = f_i \Big/ \sum_{i=1}^{n} f_i \tag{7}$$

where P_i is the probability of survival for the i^{th} genome in the population. It is clear that the higher the fitness value, the more likely it is to survive to the next generation.

It should be noted that in a gridded LPV system, the optimal states identified by the GA described above may vary across flight parameters. Therefore statistics to determine a set of common states within the entire flight regime must be undertaken to ensure state consistence. Specifically, 9 fight conditions determined by all the combinations among the lower and upper bounds and the middle point of each individual flight parameter (i.e., [50 100 150] KEAS ⊗ [0 40 80] lbs) were selected as the probes to sample the parameter space. The GA algorithm was then applied to these 9 parameter points in order to sort and determine the state to keep (or remove) for each of them. The sets of states at these conditions are then compared against each other, and those with the highest occurrences are retained and used across the entire parameter space, including those flight conditions not interrogated. In the X-56A MUTT model aerodynamic lag and elastic modes contribute to the largest number of states, and require significant experience and trial-and-error iteration to determine the optimum states to retain. Therefore, GA-guided MOR was applied to these states, and is detailed as follows:

(1) Perform GA-guided MOR at the 9 probe conditions in the parameter space, consisting of the full combination above.
(2) For the aerodynamic lag states, the number of the states to keep is tentatively set to 22. The optimal states identified by the GA approach at the 9 probes are examined to select the states with the high number of occurrence. The said state will be used in the common set of states.
(3) For the elastic states, the number of the states to keep is tentatively set to 8. The optimal states identified by the GA approach at the 9 probes are inspected to identify those with the high number of occurrence. Those states will be included in the common set of states used throughout the parameter space.

It is clear that as the requirement for the number of occurrence for a state becomes more relaxed, i.e., less occurrence, the number of states to keep in the common set will increase. For example, in the case of the full X-56A MUTT model, if the requirement is lowered from 5 to 4 and from 7 to 6 times for aerodynamic lag and elastic states, respectively, the common set of states will increase from 16 to 20 for the aerodynamic lag, and increase from 8 to 10 for the elastic states.

C Balanced Realization and Truncation for Unstable System

The presence of the parameter-dependent unstable states in the X-56A MUTT ASE model causes formidable challenges and complexities for MOR and consistent state representation. The unstable states should be retained in the reduced-order system to enable the use of ROM for synthesizing controller for stabilization and damping augmentation. Traditionally a stable/antistable separation is carried out prior to balanced truncation, and the balancing transformation is only applied to the stable parts of the system, and the unstable part remains intact (e.g., functions *balreal* and *balancmr* in Matlab). In the present effort, the balanced truncation for unstable systems based on generalized representation of the gramians proposed by Zhou [15] was developed instead. The generalized gramians of the unstable system is well-defined when there is no poles on the imaginary axis, and is given as

$$\tilde{P} := \frac{1}{2}\int_{-\infty}^{\infty} (j\omega I - A)^{-1} BB' (-j\omega I - A')^{-1} d\omega \quad \text{and} \quad \tilde{Q} := \frac{1}{2}\int_{-\infty}^{\infty} (-j\omega I - A')^{-1} C'C (j\omega I - A)^{-1} d\omega \tag{8}$$

Using a linear transformation T to achieve separation of the stable (A_s, B_s, C_s, D_s) and antistable (A_{ns}, B_{ns}, C_{ns}, D_{ns}) part of the system, the generalized controllability $\tilde{P}$ and observability $\tilde{Q}$ gramians defined in Eq. (8) can be constructed

$$\tilde{P} = T^{-1}\begin{bmatrix} P_s & 0 \\ 0 & P_{ns} \end{bmatrix}\left(T^{-1}\right)', \quad \tilde{Q} = T'\begin{bmatrix} Q_s & 0 \\ 0 & Q_{ns} \end{bmatrix}T \tag{9}$$

where P_s and Q_s are the controllability and observability gramians of (A_s, B_s, C_s), and P_{ns} and Q_{ns} are those of ($-A_{ns}$, B_{ns}, C_{ns}), which can be computed by

$$AP + PA^T + BB^T = 0 \quad \text{and} \quad A^T Q + QA + C^T C = 0 \tag{10}$$

The balancing transformation matrix for both the stable and antistable part then can be calculated as:

$$\tilde{V} = UZ\Sigma^{-1/2} \quad and \quad \tilde{W} = LY\Sigma^{-1/2} \tag{11}$$

where $\tilde{P}=UU^{\mathrm{T}}$ and $\tilde{Q}=LL^{\mathrm{T}}$ and Z, $\sum$ and Y can be obtained from singular value decomposition (SVD) $U^{\mathrm{T}}L = Z\Sigma Y^{\mathrm{T}}$. Applying the balancing transformation to the state-space model yields,

$$\begin{bmatrix} \dot{x}_r \\ y \end{bmatrix} = \begin{bmatrix} \tilde{W}^T A\tilde{V} & \tilde{W}^T B \\ C\tilde{V} & D(\rho) \end{bmatrix}\begin{bmatrix} x(t) \\ u(t) \end{bmatrix} \tag{12}$$

The state-space model in the new coordinate in Eq. (12) is balanced, and hence, its controllability and observability gramians are equal and diagonal, i.e., $\tilde{P}_r = \tilde{Q}_r = \mathrm{diag}(\sigma_1, \ldots, \sigma_r \ldots, \sigma_n)$, where $\sigma_1, \ldots, \sigma_n$ are the Hankel singular values sorted in descending order. By removing the states corresponding to low Hankel singular values (e.g., $\sigma_{r+1} \ldots \sigma_n$) and truncating their corresponding columns in $\tilde{W}$ and $\tilde{V}$, a ROM without appreciably losing important input/output energy can be obtained.

D Congruence Transformation

The aforementioned ROM steps are applied to the full-order ASE model at each grid point in the flight envelope, yielding a set of local ROMs $G_{r,i} = [A_{r,i}, B_{r,i}, C_{r,i}, D_{r,i}]$, where i denotes the i^{th} grid point in the parameter space, and G_r is the reduced system. The next step is to construct a global LPV ROM from the ASE ROM $G_{r,i}$ obtained at the grid points that can be used to obtain controlled system at arbitrary locations in the parameter space. However due to transformation matrices used in MOR (e.g., $\tilde{W}$ and $\tilde{V}$ in the balanced truncation) varying with grid points, the physical meaning of the states x_r of the individual ROMs are not consistent across the flight envelope. In other words transformed states are a weighted combination of the original states and grid-point dependent. Therefore, the ROMs cannot be interpolated directly.

To mitigate this issue, one of the most widely used methods is to project the individual ROMs onto a common subspace, followed by matrix interpolation as discussed in [3, 7, 16]. The common subspace R shared by all local ROMs is obtained by the singular value decomposition (SVD) of the concatenated right projection matrices V_i at grid points; i.e., $RS\Phi^T \approx [V_1, \ldots V_{ns}]$, where V_i is the column subspace in $\tilde{V}$ corresponding to the retained states in balanced truncation, i.e., Eq. (12). However, due to the broader parameter space and distinctly different dynamic behavior associated with the X-56A MUTT model, R of a low dimension from SVD cannot effectively capture key bases contained in all V_i, leading to poor ROM performance. Therefore in the present effort, a different approach based on congruence transformation [16] was utilized. Its principle is to set the subspace V_{n0} at a grid point n_0 as the reference, and then minimize the difference between the subspace V_i at the other grid points relative to V_{n0} through the Modal Assurance Criterion (MAC) and congruence transformation T (changes of basis). Detailed interpretation and comparison of both methods is given in [17].

A straightforward and analytical solution to computing T is to take SVD of the relative configuration matrix M constructed by the local transformation matrices V_i. The algorithm is capable of automatically detecting situations

where mode crossing and mode veering occurs. The procedure starts with selecting the projection subspace at one parameter as the reference (e.g., V_{n0}), and then iterate on computing the congruence transformation T_i at the other parameters relative to the reference [16]. In this paper the grid point at the center of the flight parameter space, i.e., (100 KEAS and 40 lbs) was selected as the reference point. The congruence transformation is summarized below:

Algorithm 1: Congruence transformation T_i at parameters i ($i \neq n_0$)
Input: Projection matrices V_i
1. Select a reference point n_0 in $i = 1, 2,$ and …, n_s and its associated subspace V_{n0}
2. For $i = 1, 2, \ldots, n_s$ and $i \neq n_0$
   Compute $M_i = V_i^T V_{n0}$
   Compute $M_i = U_i \Sigma_i Z_i^T$ via singular value decomposition (SVD)
   Compute $T_i = U_i Z_i^T$
   Compute $A^*_i = T_i^T A_{r,i} T_i$, $B^*_i = T_i^T B_{r,i}$, $C^*_i = C_{r,i} T_i$, $D^*_i = D_{r,i}$
   End for
Output: transformed system matrices G^*_i

D Interpolation for Parameter-varying Reduced Order Model

The next step is to interpolate ROMs at the grid points to obtain linear parameter-varying (LPV) ROMs that can are able to predict ASE behavior at any parameter location. At the regions where system response is susceptible to the flight parameters, the system matrices of the ROMs may change dramatically and cause significant complexities for ROM interpolation. A workaround is to add more grid points and full-scale models at these regions for refinement, guiding smooth interpolation and transition of ROMs. On the other hand a uniform refinement on the entire parameter space will give rise to a large set of models (and grid points), and require more effort for control synthesis and on-board computer storage. This is typically unnecessary as the response of the reduced system in most of the regions still varies smoothly. Therefore an adaptive refinement approach was also developed, i.e., new grid points and associated full models will be added only in the region with high interpolation error. The interpolation can be easily performed in MATLAB using a spline function. Figure 6 depicts the process of the adaptive grid refinement, which is detailed as follows:

(1) Create two maps, respectively, with 1×1 and 10×10 resolution in the 2D parameter space of KEAS and fuel weights, respectively, yielding 100×80 and 10×8 cells. They are termed the low-resolution and high-resolution map hereafter.
(2) Interpolate the full-order X-56A MUTT models on the low-resolution grids to obtain the full models on the high-resolution grids.
(3) Apply the sequential MOR approach to all the grid points in both the high- and low-resolution maps. The ROMs in the former will be used as the benchmark and ground truth, while the latter will be used to interrogate the interpolation error.
(4) Define five points in each cell of the low-resolution map as the probes representing that cell (as shown by the red dots in Figure 6).
(5) The ROMs at these five points are obtained by spline interpolation of matrix entries of ROMs on the low-resolution grid points. Compare them against their counterparts in the high resolution map, and obtain a quantitative measure for the error based on equation (5).
(6) The cells whose errors are ranked in the top 1/3 are selected for refinement, and the cell resolution therein is then doubled.
(7) In general, step 1 to 5 can be repeated till the highest resolution is reached, while in this study only one-level refinement yielding 5×5 resolution was undertaken to demonstrate the feasibility of enhancing interpolation via adaptive refinement.
(8) Save the ROMs on the low-resolution map and in the refined regions

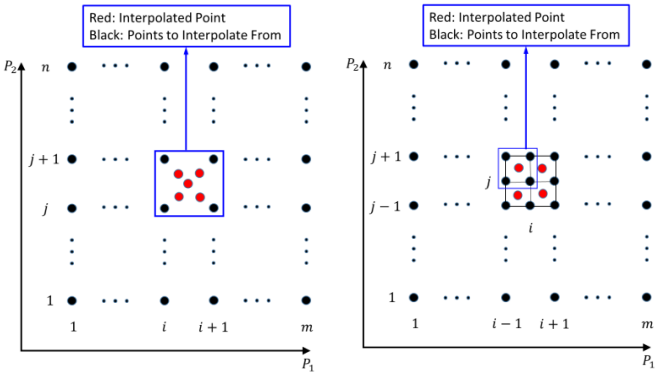

Figure 6. Diagram of interpolating an arbitrary parameter point (a) before; and (b) after grid refinement in the cell

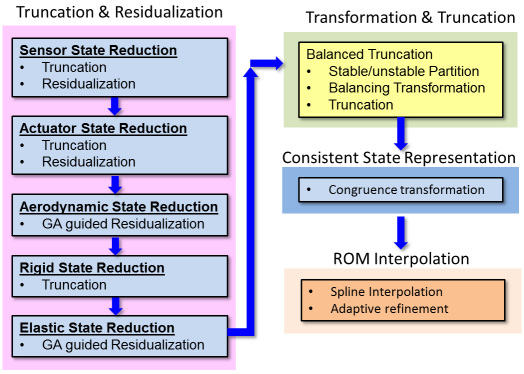

Figure 7. Process and main techniques used for MOR, consistence state representation, and ROM integration for ASE models of aircraft

V. Results and Discussion

The sequential MOR framework was applied to the X-56A MUTT ASE model described in section IV. Truncation and residualization was used to eliminate the states associated with sensors, actuators, aerodynamic states, rigid body states, and elastic states, followed by transformation-based MOR (balanced truncation for unstable systems). Changes of basis were then applied to ROMs at the grid points based on the congruence transformation to achieve the most state consistence among local ROMs, and render them ready for interpolation. The matrix entries of ROMs then can be interpolated using the spline interpolation to yield LPV ROM encompassing the entire flight regime. Figure 7 summarizes the flow chart of our MOR procedure and the main techniques being used. The investigations in this section include evaluating the effectiveness of the methodology by comparing ROMs with the full model, effects of the flight parameters on the ROM, and ROM interpolation. Due to the large number of input – output channels, only 3 input – output channels will be shown, viz., from (*WF1R* + *WF1L*)/2 to pitch rate (*qb*), and ASE 2, and from (*WF4R* – *WF4L*)/2 to roll rate (*pb*). The center point in the flight envelope (100 KEAS and 40 lbs) serves as the benchmark case.

A Sequential Model Order Reduction (MOR)

The sequential MOR framework was first performed on the benchmark case in accordance with the process as shown in Figure 7.

(1) Sensor Reduction:
Given the independent nature of the sensors states as indicated by the *A* matrix pattern in the full-order model, we first truncated 40 states in the first 56 states that have no contribution to the sensor observation. Then the remaining 16 states associated with the 1st and 2nd order sensor dynamics in study were fully residualized to match the DC gain of the original model, leading to a ROM with 136 states.

(2) Actuator Reduction:
The X-56A MUTT model includes 10 surface controls and 2 engine controls, and each is described by the 3rd-order dynamics. 12 states corresponding to four actuators that are not the object of our ASE study were truncated. Next the 3rd-order states of the remaining 8 actuators in study, i.e., WF1L– WF4L and WF1R–WF4R were residualized yielding a ROM of 116 states with the 2nd-order approximation of actuator dynamics.

(3) Aerodynamic Lag Reduction:
The X-56A MUTT model includes 60 aerodynamic lag states in the full model. In contrast to the trial and error approach in previous efforts, the GA-guided truncation and residualization was used to further reduce another 40 aerodynamic states in the 116-state model yielding a ROM of 76 states. The GA guided approach will be compared against the manual selection approach, while in the latter the first 20 leading aerodynamic lag states were retained.

(4) Rigid Body Reduction:
The X-56A MUTT model has 12 rigid body states in the full-order model, which are x, u, h, α, θ, q, y, β, p, r, ϕ, and ψ. Several comparative studies were carried out in which various combinations of the rigid-body states were examined to determine the best set that has negligible impact to system dynamics, resulting in a ROM with 68 states with consistent performance across the entire flight envelope.

(5) Elastic State Reduction:
The X-56A MUTT model has 14 coupled elastic modes, corresponding to 28 states. The first 14 states were used to describe the modal displacements, while the rest for the modal velocity. The GA-guided approach was used to keep 5 states each in modal displacement and modal velocity in the ROM, that is, 18 states in total were residualized yielding a ROM with 54 states. The GA-guided approach will again be compared against the manual selection approach, and in the latter the first 5 leading states in the modal displacement and modal velocity (10 in total) were kept.

(6) Balanced Truncation for Unstable Systems:
It was used to transform the existing ROM into a balanced controllable and observable form using the stable/antistable state separation and generalized balancing transformation. The states in the model were sorted according to the significance of their corresponding Hankel singular values. Therefore, the states with the least controllability and observability were truncated to construct the minimal realization of the model capable of capturing the dynamics between all input-output pairs. In this study, 27 states with smaller Hankel singular values were truncated, yielding a ROM with 27 states.

The sequential MOR and resulting model sizes are summarized in Table 3. Figure 8 through Figure 13 show the comparison of the magnitude and phase in the frequency domain between the full-order X-56A and ROM for the benchmark case (100 KEAS 40 lbs) during the sequential MOR. It demonstrates that ROM accurately matches the full-order model for all the selected input-output channels within the desired frequency range while the number of states is reduced by almost 7X.

The GA-ROM and manual-ROM are exactly the same in Figure 8 and Figure 9 as the GA-guided state selection and reduction has not been applied. The difference between them is clearly observed in Figure 10, where GA was used to automatically determine the states to remove with the least compromise to the original system dynamics while in the manual selection the first 20 leading states are retained. The GA-ROM outperforms the manual ROM dramatically in particular in the channels of pitch and roll control, e.g., from (WF1L+WF1R)/2 to *qb* and from (WF4L-WF4R)/2 to *pb*. The salient performance was carried over to the downstream analysis. In particular in the elastic (flexible) state reduction (Figure 12), GA-ROM and manual ROM were used to only keep 10 states (18 states residualized). The former still resolved the dynamics among all channels very well, while the latter failed to capture the system response at the high-frequency regime in the roll channel. Figure 13 illustrates the comparison of the ROM obtained by the balanced truncation against the original model. The model of 54 states from elastic reduction was transformed into the balanced form, and 27 states with smallest Hankel values were then truncated, resulting in a final ROM with only 27 states.

Table 3. Sequential MOR and resulting model sizes

Reduction	Original	Sensor	Actuator	Aerodynamic	Rigid-Body	Elastic	Balanced Truncation
Model Size	192	136	116	76	72	54	27

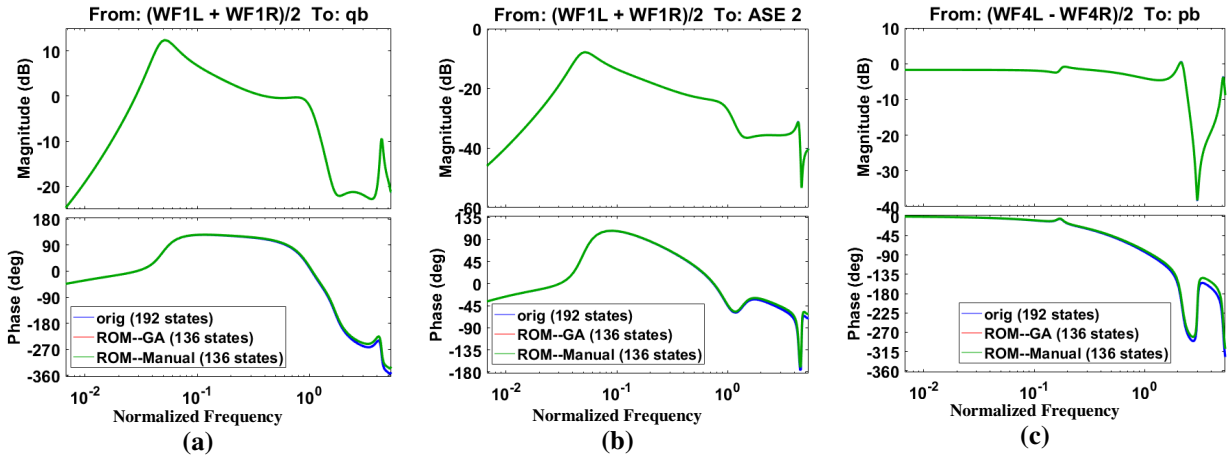

Figure 8: Comparison in magnitude and phase in the frequency domain between the full-order X-56A model (192 states) and the ROM after sensor reduction. Pitch control: From (WF1L+WF1R)/2 to (a) qb and (b) ASE 2; Roll control: (c) From (WF4L-WF4R)/2 to pb.

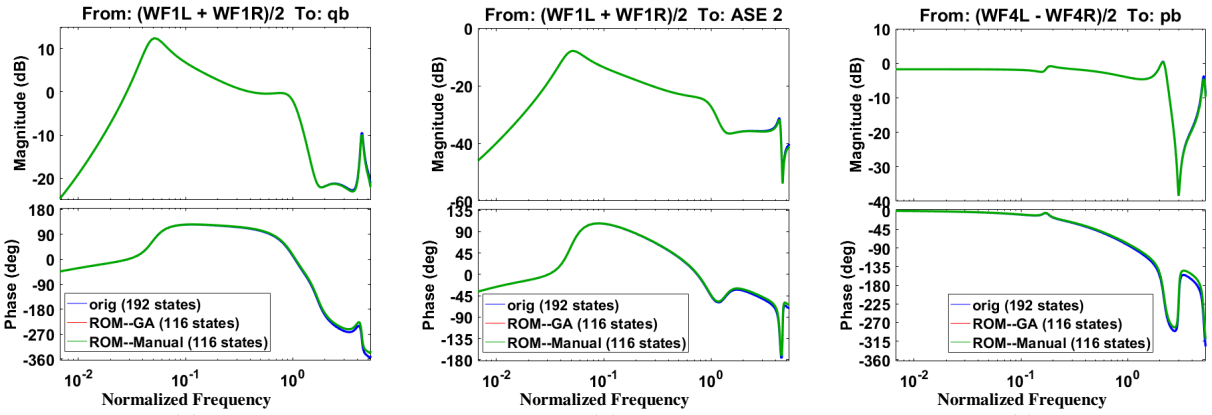

Figure 9: Comparison in magnitude and phase in the frequency domain between the ROM before and after actuator reduction. Pitch control: From (WF1L+WF1R)/2 to (a) qb and (b) ASE 2; Roll control: (c) From (WF4L-WF4R)/2 to pb.

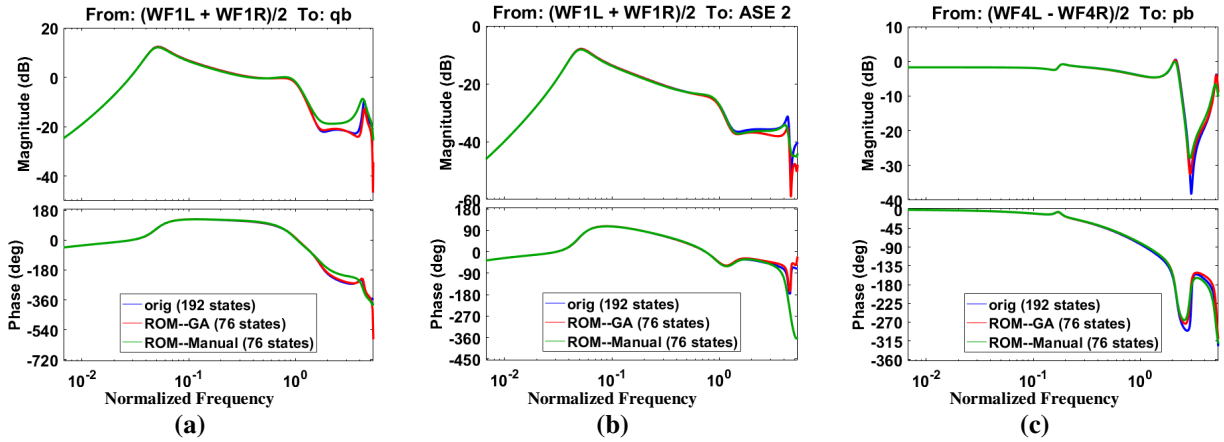

Figure 10: Comparison in magnitude and phase in the frequency domain between the ROM before and after aerodynamic lag reduction. Pitch control: From (WF1L+WF1R)/2 to (a) qb and (b) ASE 2; Roll control: (c) From (WF4L-WF4R)/2 to pb.

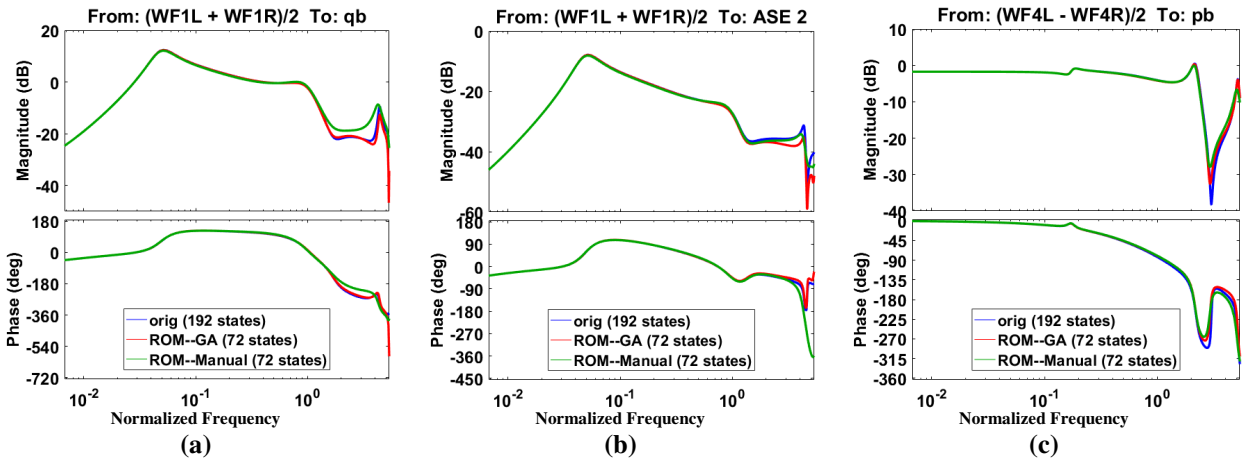

Figure 11: Comparison in magnitude and phase in the frequency domain between the ROM before and after rigid-body reduction. Pitch control: From (WF1L+WF1R)/2 to (a) qb and (b) ASE 2; Roll control: (c) From (WF4L-WF4R)/2 to pb.

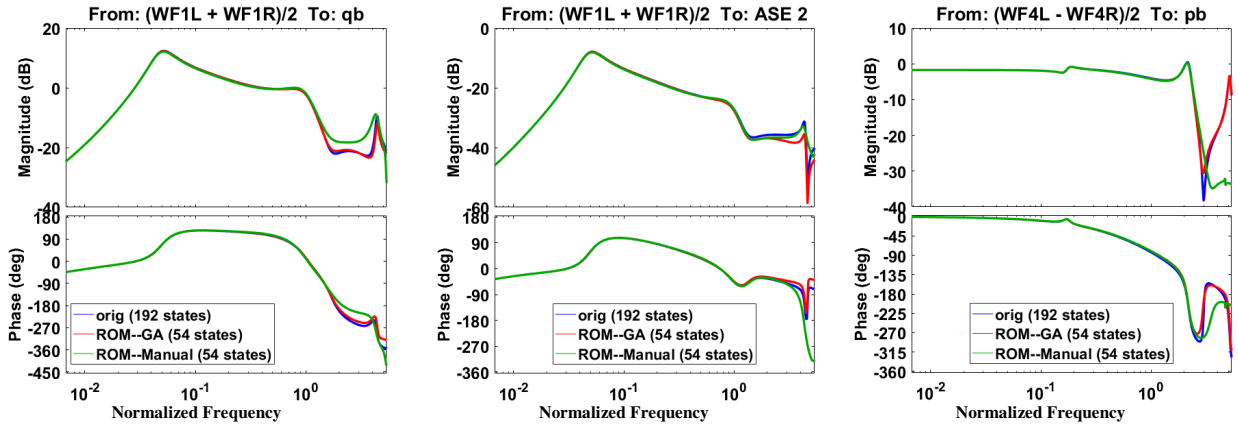

Figure 12: Comparison in magnitude and phase in the frequency domain between the ROM before and after elastic state reduction. Pitch control: From (WF1L+WF1R)/2 to (a) qb and (b) ASE 2; Roll control: (c) From (WF4L-WF4R)/2 to pb.

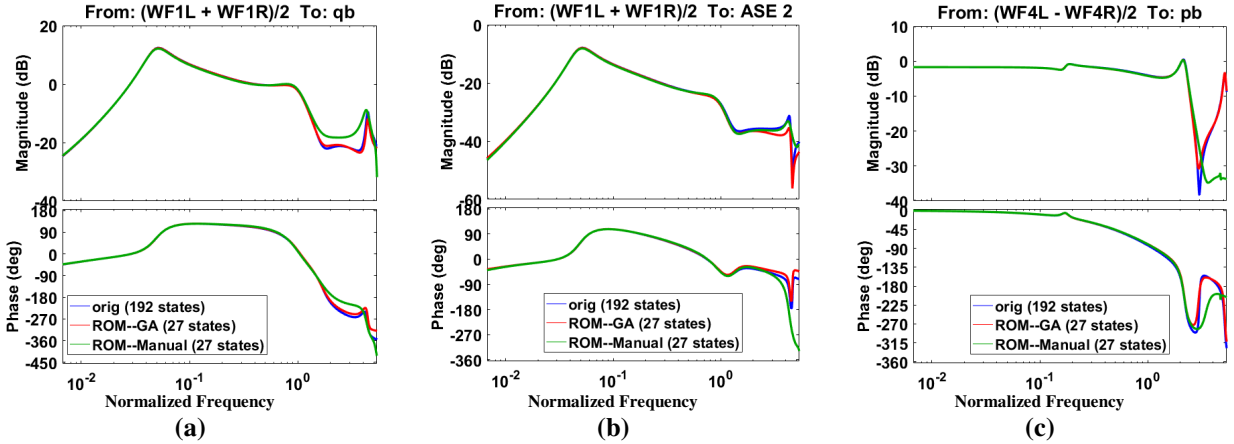

Figure 13: Comparison in magnitude and phase in the frequency domain between the ROM before and after balanced truncation. Pitch control: From (WF1L+WF1R)/2 to (a) qb and (b) ASE 2; Roll control: (c) From (WF4L-WF4R)/2 to pb.

B ROM with Different Flight Parameters

A desired feature that will saliently enhance the utility of MOR for aircraft ASE analysis and controller synthesis is its robustness and consistence of model configuration parameters regardless of the flight conditions. This is the exact reason that the GA-guided MOR has been performed at the representative parameters in the domain to identify the common aerodynamic and elastic states for reduction among all ROMs through statistics (see Section IV. B). Therefore a thorough study to investigate the effect of varying flight parameters (KEAS and fuel weights) on MOR performance was also carried out in the present effort. Figure 14 shows the comparison between the original 192-state, GA-ROM, and manual-ROM at 4 flight parameters (80 KEAS 20 lbs, 80 KEAS 70lbs, 120 KEAS 20 lbs, 120 KEAS 70 lbs) different from the benchmark case. Note that all the configuration parameters for the serial MOR remained the same as the case study in Section A above. The GA-ROMs at the new flight conditions are able to accurately describe the dynamics in all input-output channels (pitch, roll, and ASE) in the frequency range of interest. The ROMs deviate from the original 192-state models only at the middle-to-high frequency regime, verifying pronounced robustness and utility of our MOR methods.

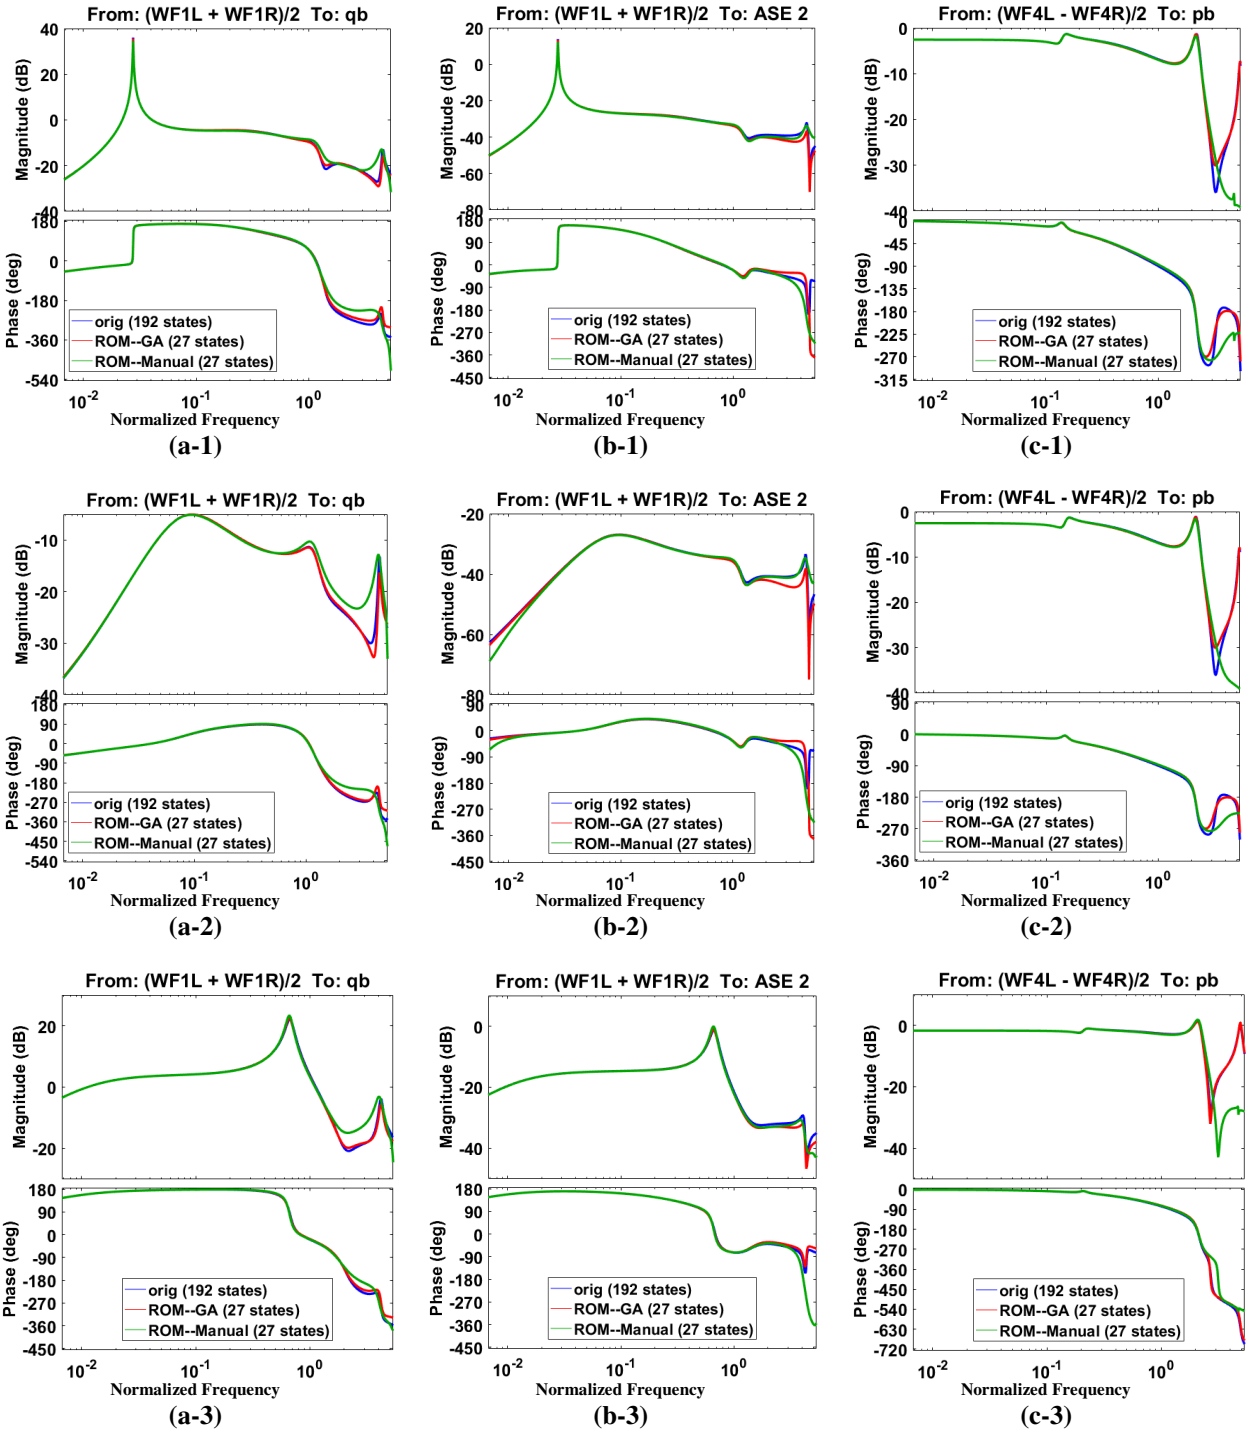

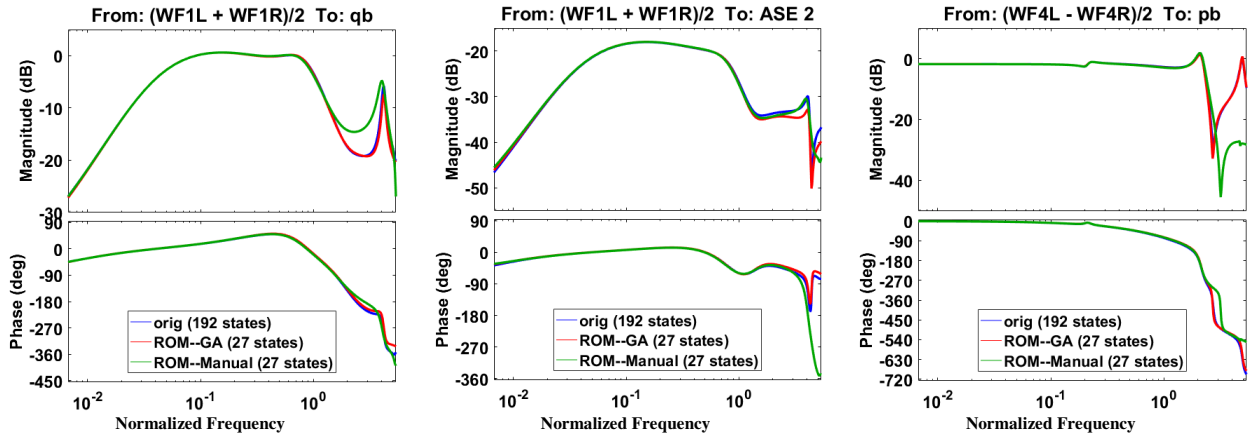

(a-4) (b-4) (c-4)

Figure 14: Comparison in magnitude and phase in the frequency domain between the ROM before and after balanced truncation. Pitch control: From (WF1L+WF1R)/2 to (a) qb and (b) ASE 2; Roll control: (c) From (WF4L-WF4R)/2 to pb. Row 1: 80 KEAS 20 lbs, Row 2: 80 KEAS 70lbs, Row 3: 120 KEAS 20 lbs, Row 4: 120 KEAS 70 lbs

C ROM Interpolation

The consistent state representation and ROM interpolation to achieve LPV ASE ROMs based on congruence transformation, adaptive grid refinement, and ROM interpolation as formulated above was also investigated.

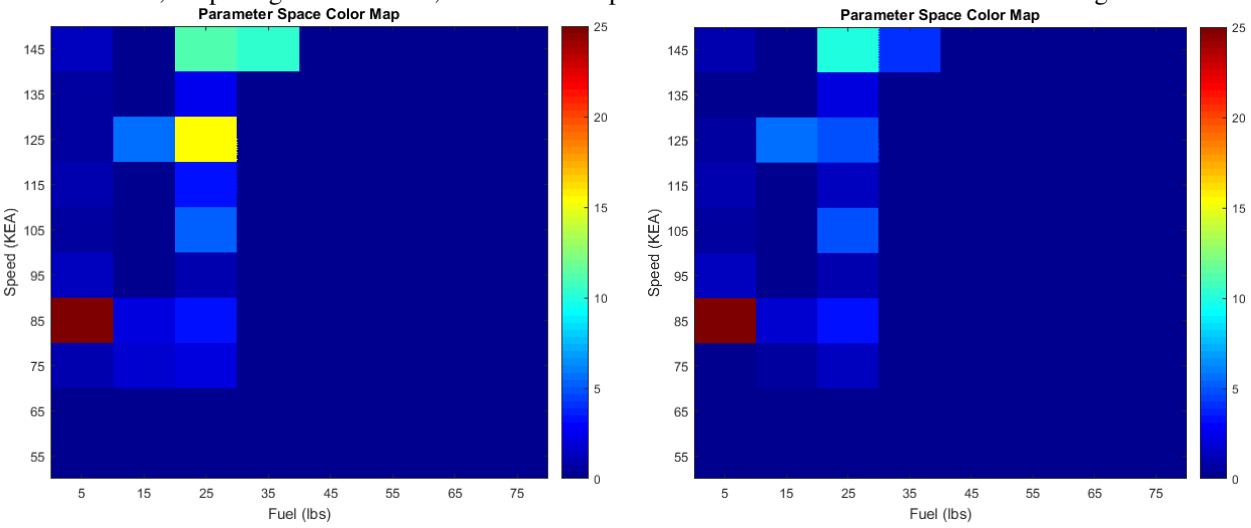

(a) (b)

Figure 15: Color Map indicating the interpolation error (a) on the low-resolution map and before adaptive refinement; and (b) after adaptive refinement of the grids where interpolation error.

Figure 15a shows the color map indicating the error of the ROM interpolation (based on Eq. (5)) on the low-resolution grid (i.e., 10 KEAS × 10 lbs). 1/3 region with the largest errors in the parameter space as illustrated by the blocks in the light color were then refined according to the procedure described in section IV-D. Figure 15b illustrates the color map of interpolation error after adaptive refinement (i.e., 5 KEAS × 5 lbs). The errors in the refined regions all drop down, but to different extents. For example the yellow block covering the flight parameter of 127 KEAS and 23 lbs exhibits significant reduction in error, which is confirmed by comparing the "Interpolated System" before and after the adaptive refinement against the one without interpolation ("True System") as shown in Figure 16. On the other hand the interpolated ROM only improves slightly at the flight parameter of (83 KEAS and 7 lbs). It can be attributed to the fact that the H-infinitely norm was used as a measure of interpolation error in our adaptive refinement process. While the maximum error was not notably reduced at the peak response (see the bottom row in Figure 17), the overall enhancement in interpolation accuracy is evident across the entire frequency

range. Therefore another refinement at the next level of grid resolution may be exploited for additional accuracy. Although H_2-norm is not available for the unstable system, such as our X-56A MUTT model, other definitions suited for objective quantification of interpolation errors across the entire frequency range will be considered in the future (e.g., probed at multiple, representative frequencies).

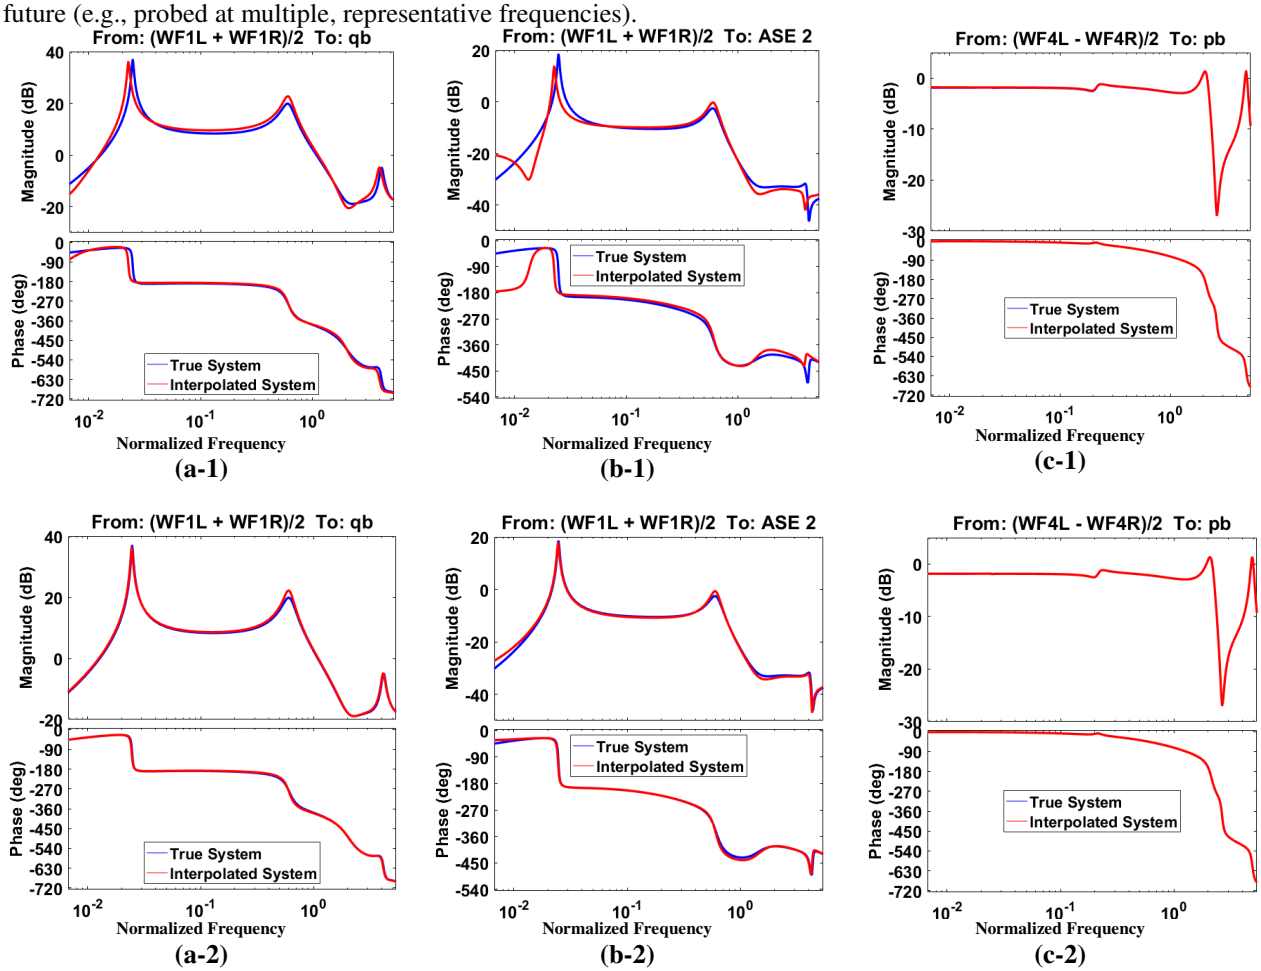

Figure 16: Comparison in magnitude and phase in the frequency domain between the ROM before and after adaptive refinement. Pitch control: From (WF1L+WF1R)/2 to (a) qb and (b) ASE 2; Roll control: (c) From (WF4L-WF4R)/2 to pb. Interpolation with resolution of (Top Row) 10 KEAS × 10 lbs and (Bottom Row) 5 KEAS × 5 lbs at 127 KEAS and 23 lbs.

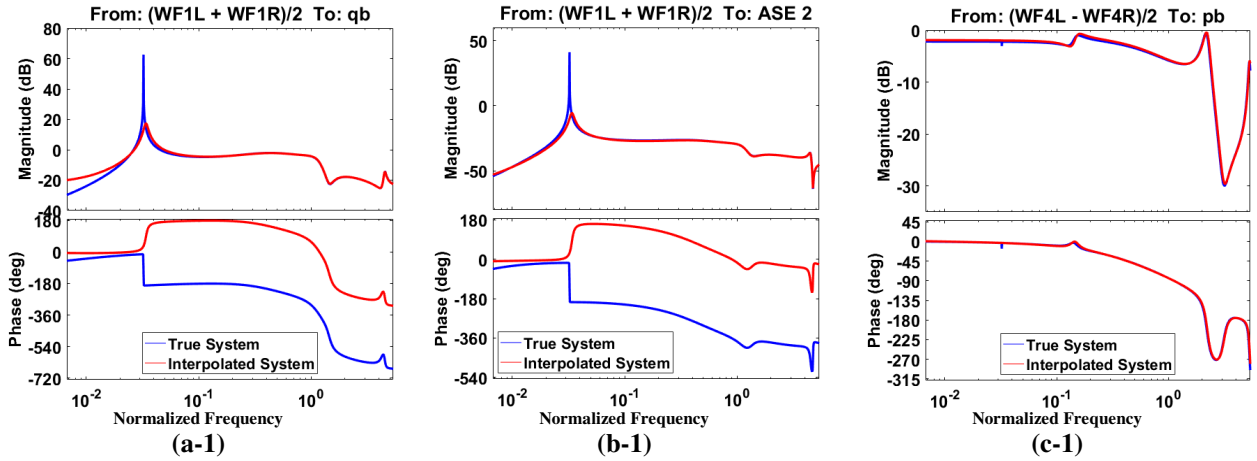

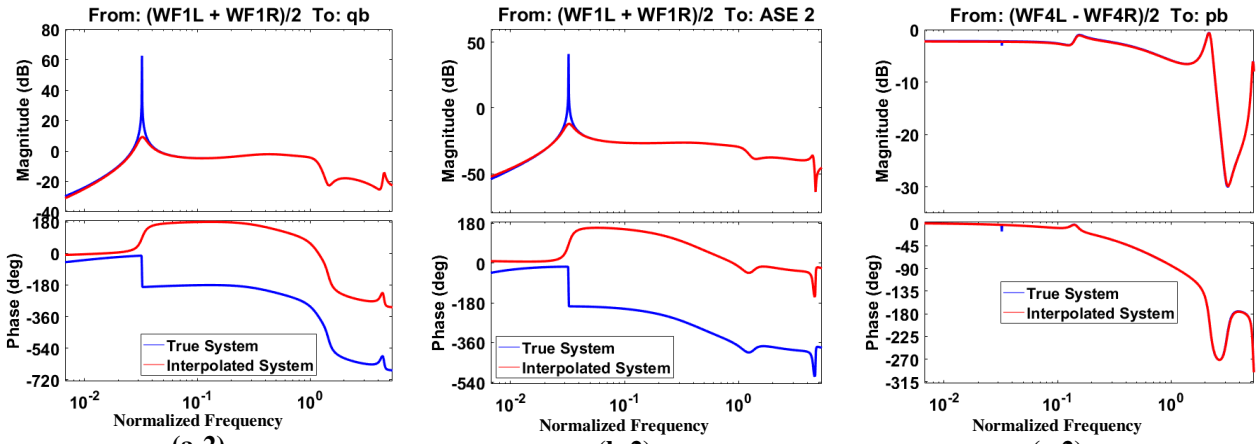

Figure 17: Comparison in magnitude and phase in the frequency domain between the ROM before and after adaptive refinement. Pitch control: From (WF1L+WF1R)/2 to (a) qb and (b) ASE 2; Roll control: (c) From (WF4L-WF4R)/2 to pb. Interpolation with resolution of (Top Row) 10 KEAS × 10 lbs and (Bottom Row) 5 KEAS × 5 lbs at 83 KEAS and 7 lbs.

VI. Conclusion

This paper presents a systematic framework for sequential model order reduction (MOR) of high-dimensional, unstable, linear parameter varying aeroservoelastic (LPV ASE) models of flexible aircrafts. Key MOR components including sequential model order reduction, consistent model representation, and model interpolation based on adaptive refinements have been developed to establish feasible workflow. The novelty of the present effort includes the genetic algorithm (GA) for automated state selection for reduction; the balanced truncation for unstable systems to further condense model sizes; use of the congruence transformation to unify the state representation across the entire 2D flight parameter space, and the adaptive grid refinement and ROM interpolation to construct a globally valid LPV ASE ROM. The MOR technology was verified by the NASA X-56A MUTT model that includes flexible wing aerodynamics, sensors, actuators, etc. Our studies demonstrate that the ROM featuring >7X state reduction (from 192 states to 27 states) was capable of accurately capturing aircraft dynamics and system response among all relevant input-output channels (pitch, roll, and ASE) in the practically important frequency regime at various flight conditions. The GA-guided approach saliently outperforms the manual and empirical selection. The adaptive refinement allows selective addition of the grid points in the broad flight parameter space to markedly enhance the interpolation accuracy while minimizing the overhead of controller synthesis and requirements of onboard memory downstream. The technology enables robust and efficient ASE controller synthesis for aircraft, novel vehicle design for flutter suppression and gust load alleviation, notable reduction in development time and cost.

Acknowledgments

This research is sponsored by NASA under contract NNX15CD03C.

References

1. Hjartarson, A., P.J. Seiler, and G.J. Balas. *Lpv aeroservoelastic control using the lpvtools toolbox*. in *AIAA Atmospheric Flight Mechanics (AFM) Conference*. 2013.
2. Moreno, C.P., P.J. Seiler, and G.J. Balas, *Model Reduction for Aeroservoelastic Systems.* Journal of Aircraft, 2014. **51**(1): p. 280-290.
3. Panzer, H., et al., *Parametric model order reduction by matrix interpolation.* at-Automatisierungstechnik Methoden und Anwendungen der Steuerungs-, Regelungs-und Informationstechnik, 2010. **58**(8): p. 475-484.
4. Poussot-Vassal, C. and C. Roos. *Flexible aircraft reduced-order LPV model generation from a set of large-scale LTI models*. in *American Control Conference (ACC), 2011*. 2011. IEEE.
5. Poussot-Vassal, C. and F. Demourant, *Dynamical medium (large)-scale model reduction and interpolation with application to aircraft Systems.* AerospaceLab, 2012(4): p. p. 1-11.
6. Theis, J., et al. *Modal Matching for LPV Model Reduction of Aeroservoelastic Vehicles*. in *AIAA Atmospheric Flight Mechanics Conference*. 2015. Kissimmee, Florida.

7. Wang, Y., et al. *Model Order Reduction of Aeroservoelastic Model of Flexible Aircraft*. in *57th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, AIAA SciTech*. 2016.
8. *http://www.mscsoftware.com/product/msc-nastran*.
9. *http://www.zonatech.com/ZAERO.htm*.
10. Pak, C.-g. and S. Truong. *Creating a Test Validated Structural Dynamic Finite Element Model of the X-56A Aircraft*. in *5TH AIAA/ISSMO MULTIDISCIPLINARY ANALYSIS AND OPTIMIZATION CONFERENCE*. 2014.
11. Moreno, C., et al. *Model Reduction of Flexible Aircraft for Flutter Suppression using Smart Sensors*. in *Proc. of the 6th International Workshop on Advanced Smart Materials and Smart Structures Technology, Dalian, China*. 2011.
12. Mitchell, M., *An Introduction to Genetic Algorithms*. 1996: MIT Press.
13. Beasley, D., D.R. Bull, and R.R. Martin, *An overview of genetic algorithms: Part 2, research topics.* University computing, 1993. **15**(4): p. 170-181.
14. Beasley, D., R. Martin, and D. Bull, *An overview of genetic algorithms: Part 1. Fundamentals.* University computing, 1993. **15**: p. 58-58.
15. Zhou, K., G. Salomon, and E. Wu, *Balanced realization and model reduction for unstable systems.* International Journal of Robust and Nonlinear Control, 1999. **9**(3): p. 183-198.
16. Amsallem, D. and C. Farhat, *An online method for interpolating linear parametric reduced-order models.* SIAM Journal on Scientific Computing, 2011. **33**(5): p. 2169-2198.
17. Geuss, M., H. Panzer, and B. Lohmann. *On parametric model order reduction by matrix interpolation*. in *Proceedings of the 2013 European Control Conference (ECC)*. 2013. Zürich, Switzerland.