Interface reduction for Hurty/Craig-Bampton substructured models

Review and improvements

Krattiger, Dimitri; Wu, Long; Zacharczuk, Martin; Buck, Martin; Kuether, Robert J.; Allen, Matthew S.; Tiso, Paolo; Brake, Matthew R.W.

DOI
10.1016/j.ymssp.2018.05.031

Publication date
2019

Document Version
Accepted author manuscript

Published in
Mechanical Systems and Signal Processing

Interface Reduction for Hurty/Craig-Bampton Substructured Models: Review and Improvements

Dimitri Krattiger[a], Long Wu[b], Martin Zacharczuk[c], Martin Buck[c], Robert J. Kuether[d], Matthew S. Allen[e], Paolo Tiso[f], Matthew R. W. Brake[g]

[a]*Department of Aerospace Engineering Sciences, University of Colorado Boulder, Boulder, Colorado 80309, USA*
[b]*Department of Maritime and Materials Engineering, Delft University of Technology, 2628CD Delft, The Netherlands*
[c]*University of Stuttgart, 70569 Stuttgart, Germany*
[d]*Sandia National Laboratories, Albuquerque, New Mexico 87185, USA*
[e]*Department of Engineering Physics, University of Wisconsin-Madison, Madison, Wisconsin 53706, USA*
[f]*Institute for Mechanical Systems, ETH Zürich, 8092 Zürich, Switzerland*
[g]*Department of Mechanical Engineering, William Marsh Rice University, Houston, Texas 77005, USA*

Abstract

The Hurty/Craig-Bampton method in structural dynamics represents the interior dynamics of each subcomponent in a substructured system with a truncated set of normal modes and retains all of the physical degrees of freedom at the substructure interfaces. This makes the assembly of substructures into a reduced-order system model relatively simple, but means that the reduced-order assembly will have as many interface degrees of freedom as the full model. When the full-model mesh is highly refined, and/or when the system is divided into many subcomponents, this can lead to an unacceptably large system of equations of motion. To overcome this, interface reduction methods aim to reduce the size of the Hurty/Craig-Bampton model by reducing the number of interface degrees of freedom. This research presents a survey of interface reduction methods for Hurty/Craig-Bampton models, and proposes improvements and generalizations to some of the methods. Some of these interface reductions operate on the assembled system-level matrices while others perform reduction locally by considering the uncoupled substructures. The advantages and disadvantages of these methods are highlighted and assessed through comparisons of results obtained from a variety of representative linear FE models.

Keywords: Component Mode Synthesis, Substructuring, Hurty/Craig-Bampton Method, Interface Reduction, Characteristic Constraint Modes

1. Introduction

Large and complex structures often consist of multiple components and therefore can be decomposed into an assembly of relatively simple subsystems. Component mode synthesis (CMS), also referred to as dynamic substructuring, is widely used to assemble the reduced models of substructures and form a global system. Comprehensive reviews of CMS methods are found in [6, 7, 16]. Substructuring techniques were applied to multi-component systems as early as the 1960s by Hurty [13]. He also proposed the first CMS technique based on fixed-interface modes, rigid body modes, and constraint modes (CMs) [14]. Hurty's approach was simplified later by Craig and Bampton who realized that the rigid body modes need not be considered separately if all interface degrees of freedom (DOF) are included in the CMs [4]. To this day, the Hurty/Craig-Bampton (HCB) approach is one of the most popular CMS techniques in industry and academia. The HCB transformation matrix uses a truncated set of fixed-interface modes and a set of constraint modes to reduce the partitioned system of equations. Each constraint mode shape is the static deformation of the structure obtained by applying a unit displacement to one interface DOF while all the other interface DOF are held fixed. The fixed-interface mode shapes are computed from an eigenvalue analysis while fixing each interface DOF of the subcomponent. Another class of CMS methods uses free-interface modes to represent the internal dynamics of the substructures and attachment modes to represent the interface dynamics [20, 23, 8]. More recently, Rixen [22] proposed a dual Craig-Bampton method where weak interface compatibility is used to avoid

interface locking when the reduction basis contains very few modes. Weak compatibility refers to an interface coupling strategy that allows for non-matching bases at the interfaces, and allows for (typically small) discontinuities in displacement across the interfaces. Requiring the discontinuities to be too small (or zero) results in an overly stiff or locked interface. CMS techniques work well for finding the low frequency modes of a system but do not perform well for mid- to high-frequency modes. A few methods have been developed to extend CMS to higher frequencies by replacing the static CMs with quasi-static CMs and by applying shifting around a central frequency [18, 25, 24].

CMS techniques commonly focus on reduction of the interior DOF of each substructure but retain all physical DOF at the interfaces, making the assembly straightforward. In many cases, the size of the CMS model is dominated by the interface DOF, for example when the finite-element (FE) mesh is fine and/or the number of subcomponents is large. Efficient interface reduction techniques have been developed to address this issue. In 1977, Craig and Chang [9] first proposed three methods to reduce the interface DOF by applying either Guyan, Ritz, or modal reduction. However, these methods do not appear to have been widely adopted as they are neither subsequently mentioned in the literature nor are they available in commercial codes that implement the HCB method. Castanier et al. [5] rediscovered the modal method by applying a secondary eigenvalue analysis on the interface partition of the assembled CMS model to obtain the characteristic constraint (CC) modes. The use of a truncated set of CC modes results in a highly reduced system model, however these modes must be calculated only after the substructures have been assembled and the system-level HCB matrices are available. For this reason, the method is referred to as the system-level characteristic constraint (S-CC) mode method. The need to assemble the system before reduction limits the versatility of the S-CC method and has inspired the investigation of other interface reduction methods that perform the reduction before assembling the substructure models.

Hong et al. [12] proposed an interface reduction technique that performs an eigenvalue analysis on the HCB interface DOF at the substructure level prior to assembly. They enforce exact interface compatibility by concatenating these local-level characteristic constraint (L-CC) modes from each interface and use this augmented mode set to reduce the number of DOF at each of the interfaces. Alternatively, Kuether et al. [17] proposed a weak interface compatibility method where the L-CC modes of each substructure remain separate but are linearly combined to minimize the compatibility error between connecting substructure interfaces. Lindberg et al. [19] proposed a CMS approach using undeformed coupling interface modes which constrain the interfaces to three displacement and three rotational DOF. This technique is valid for specific systems where the connection occurs between a soft and stiff subcomponent such that the deformation at the interface is governed by the stiffer structure. Aoyama et al. [2] presented a method similar to S-CC except that each interface is assembled and then reduced separately. The hybrid-level characteristic constraint (H-CC) mode method presented subsequently is an extension of this. Using their approach, the modes used at every interface depend only on the adjacent substructures and can be quickly updated if substructures are added, removed, or changed. Holzwarth et al. [11] recently demonstrated the use of Legendre Polynomials as basis functions for local interface reduction and compared the performance to S-CC interface reduction. The work by Balmès [3] explored a CMS basis that defined arbitrary interface deformations to describe a set of generalized DOF along an interface.

Due to its popularity for CMS problems, most interface reduction approaches are developed for HCB models. Tran [26, 27] showed that S-CC interface reduction can be extended to free and hybrid interface methods by first describing the attachment modes in terms of CMs and then proceeding as before. Rixen [21] proposed an interface reduction approach for the Dual Craig Bampton method, where the interface force modes are applied based on the dual nature of the connection between substructures. Interface reduction methods have also been applied to fluid-structure coupling systems using appropriate Ritz vectors [10].

The goal of this paper is to provide an overview of a variety of interface reduction techniques as applied to HCB substructure models coupled using primal assembly. The six different interface reduction strategies that are considered in this paper are the S-CC method [5], the exact-compatibility (EC) [12], weak-compatibility (WC), and uncoupled weak-compatibility (UWC) L-CC methods, the virtual node (VN) method [19], and a hybrid method that is dubbed the H-CC method. The WC L-CC and UWC L-CC techniques presented here are both generalizations of the approach presented for single interface systems in [17] to accommodate systems with multiple interfaces. The H-CC method is a modification of the method presented in [2] to allow for the calculation of rigid-body modes when a system is free. The VN method generalizes the approach in [19] to allow for multiple virtual nodes per interface. The schematic in Figure 1 classifies the various interface reduction techniques as either system-level, local-level, or hybrid-level. The phrase "system-level" indicates that interface reduction is applied to the fully assembled HCB model, whereas "local-level" indicates that interface reduction calculations are performed on the individual substructures before they

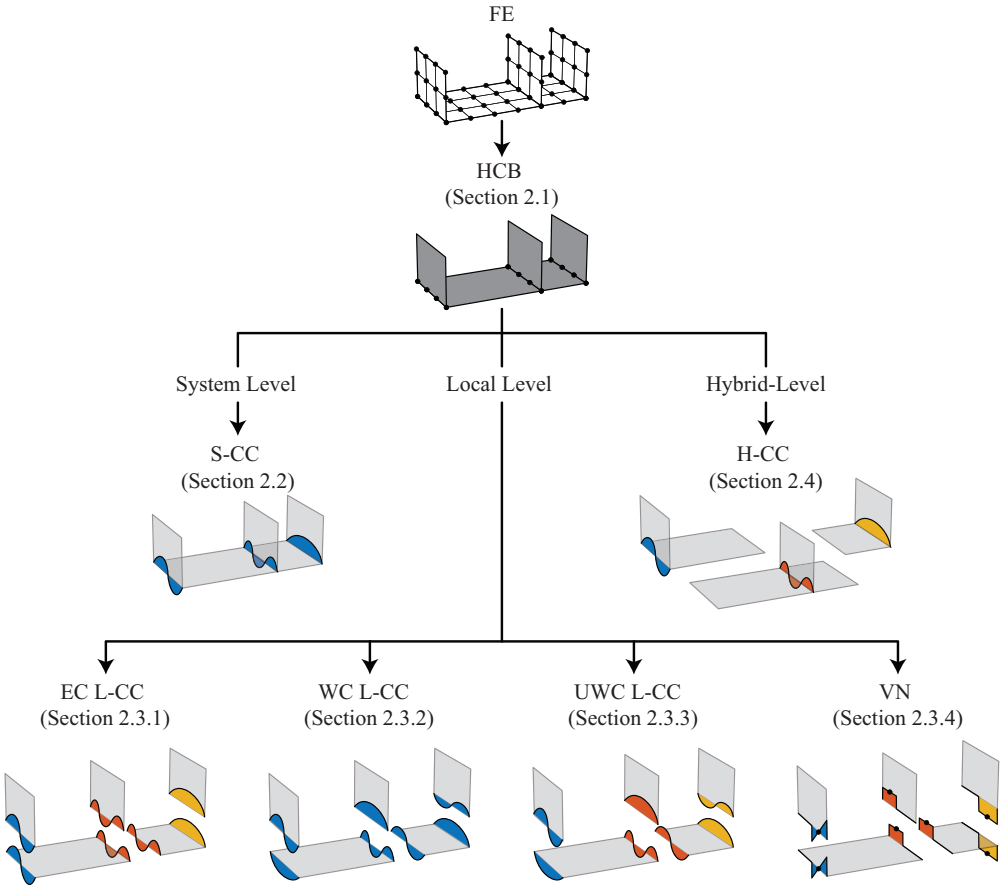

Figure 1: Schematic of the different interface reduction methods considered in this paper, and their classification as system-, local-, or hybrid-level. The thumbnail illustrations show whether interface modes are coupled (single color) or uncoupled (different colors), the level of substructure assembly when the interface modes are computed, and for the local methods, whether modes match at the interface (indicating exact compatibility) or not (indicating weak compatibility).

are assembled. "Hybrid-level" falls somewhere in between because only the substructures connecting to a particular interface need to be assembled to compute the reduction of that individual interface.

As all of the cited methods were compared, similarities and differences were noted that led to the new methods, generalizations and improvements that are presented in this paper. Furthermore, every effort will be taken in the following to clarify the exposition of the methods, highlight the similarities and differences, and explore the effect of the parameters that the user must define such as, for instance, singular value tolerances or the number of interface modes to be retained. Even then, the interface reduction method that best suits a specific engineering analysis will depend on the problem of interest, and one of the goals of this paper is to help the reader better understand the scenarios in which each method works best. The methods are investigated with linear example systems, but one of the motivations of this work is the eventual extension to model-order reduction for nonlinear systems. Section 2 of this paper reviews the classic HCB method and presents the theory for the six different interface reduction techniques. In Section 3, the performance of reduced interface models are demonstrated using three numerical examples. Conclusions are presented in Section 4.

2. Interface Reduction Techniques

This section presents the formulation of the six different interface reduction strategies that are considered in this paper. As the methods are all applied to HCB reduced-order models, a brief discussion of the theory is also included.

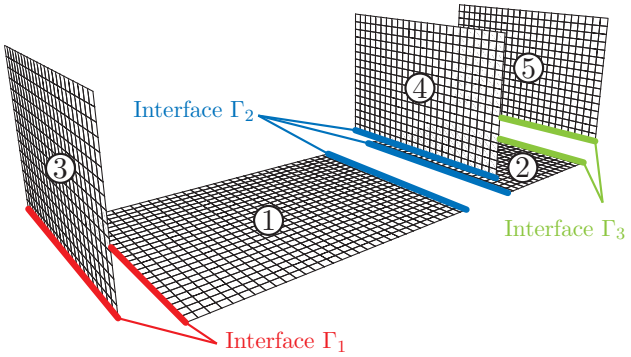

Figure 2: Exploded view of the W-bracket system, inspired by the model in [12], consisting of five plate substructures arranged in a "W" configuration with three interface sets; the interface sets Γ_i, $i = 1, 2, 3$ are highlighted and labeled.

Throughout this section, the model in Fig. 2 is used to support the theoretical discussion of each interface reduction method.

2.1. Hurty/Craig-Bampton Method

The FE discretized system of equations of motion for a substructure reads

$$\mathbf{M}_j\ddot{\mathbf{u}}_j + \mathbf{K}_j\mathbf{u}_j = \mathbf{f}_j, \tag{1}$$

where $\mathbf{M}_j$ and $\mathbf{K}_j$ are the mass and stiffness matrices, respectively, $\mathbf{u}_j$ is the displacement vector, $\mathbf{f}_j$ is the forcing vector and the subscript j denotes the j^{th} substructure. The substructure equations of motion are partitioned into interior and interface DOF (denoted by subscripts i and b respectively), as

$$\begin{bmatrix} \mathbf{M}_{ii,j} & \mathbf{M}_{ib,j} \\ \mathbf{M}_{bi,j} & \mathbf{M}_{bb,j} \end{bmatrix} \begin{Bmatrix} \ddot{\mathbf{u}}_{i,j} \\ \ddot{\mathbf{u}}_{b,j} \end{Bmatrix} + \begin{bmatrix} \mathbf{K}_{ii,j} & \mathbf{K}_{ib,j} \\ \mathbf{K}_{bi,j} & \mathbf{K}_{bb,j} \end{bmatrix} \begin{Bmatrix} \mathbf{u}_{i,j} \\ \mathbf{u}_{b,j} \end{Bmatrix} = \begin{Bmatrix} \mathbf{0} \\ \mathbf{f}_{b,j} \end{Bmatrix} \tag{2}$$

where the subscripts b and i are indexes referring to the boundary and internal component of the matrices and vectors.

The CMs are formed by computing the static response of the interior of the substructure when one interface DOF is given a unit displacement and all the other DOF are held fixed. The CM set for the entire interface is expressed as

$$\mathbf{\Psi}_j = \begin{bmatrix} -\mathbf{K}_{ii,j}^{-1}\mathbf{K}_{ib,j} \\ \mathbf{I} \end{bmatrix}. \tag{3}$$

Static condensation is a technique that removes a subset of the DOF that experience no external forcing from a model. This is accomplished by approximating the displacement of the removed set with the static response that those DOF experience due to displacement of the retained DOF set. The statically reduced model exactly reproduces the full model solutions when the system experiences only static deformation, or if the removed DOF have zero mass. Since mass terms are rarely zero, the accuracy of a statically reduced model quickly deteriorates as the frequency of vibration increases. The CMs shown above provide a basis that can be used to perform a static condensation that retains all interface DOF and eliminates all interior DOF from the model. This is powerful when applied to a substructure because the original physical interface DOF are retained, so the system can still be reassembled directly. The CMs alone cannot accurately capture the dynamic behavior of the internal DOF. This is addressed by augmenting the CMs with a set of dynamic modes that are obtained by fixing the interface DOF and assuming a harmonic solution as

$$\left(\mathbf{K}_{ii,j} - \omega_r^2\mathbf{M}_{ii,j}\right)\left\{\phi_{i,j}\right\}_r = 0. \tag{4}$$

The eigenvectors obtained from this equation are referred to as fixed-interface modes. A truncated set of these mass normalized eigenvectors are collected into a fixed-interface mode matrix,

$$\mathbf{\Phi}_j = \begin{bmatrix} \left\{\phi_{i,j}\right\}_1, \ldots, \left\{\phi_{i,j}\right\}_n \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{\Phi}_{i,j} \\ \mathbf{0} \end{bmatrix}. \tag{5}$$

These modes provide a normal basis for the interior DOF of the substructure. The fixed-interface modes and the CMs are combined to form the HCB reduction matrix as

$$\mathbf{T}_j^{\text{HCB}} = \begin{bmatrix} \boldsymbol{\Phi}_j & \boldsymbol{\Psi}_j \end{bmatrix}, \tag{6}$$

which provides a transformation from the HCB generalized DOF to the substructure physical DOF,

$$\begin{Bmatrix} \mathbf{u}_{i,j} \\ \mathbf{u}_{b,j} \end{Bmatrix} \approx \mathbf{T}_j^{\text{HCB}} \begin{Bmatrix} \mathbf{q}_{i,j} \\ \mathbf{u}_{b,j} \end{Bmatrix}, \tag{7}$$

where $\mathbf{q}_{i,j}$ represents the modal coordinate vector associated with the fixed-interface modes. The uncoupled substructure reduced mass and stiffness matrices are now formed by applying the HCB transformation to these matrices as

$$\mathbf{M}_j^{\text{HCB}} = \left(\mathbf{T}_j^{\text{HCB}}\right)^{\text{T}} \mathbf{M}_j \mathbf{T}_j^{\text{HCB}}, \qquad \mathbf{K}_j^{\text{HCB}} = \left(\mathbf{T}_j^{\text{HCB}}\right)^{\text{T}} \mathbf{K}_j \mathbf{T}_j^{\text{HCB}}. \tag{8}$$

The HCB component-level reduced-order models are coupled using a primal assembly by defining a transformation between the uncoupled and coupled DOF to eliminate the constraint equations. For the W-bracket system in Fig. 2, the coupling matrix is built directly as

$$\begin{Bmatrix} \mathbf{q}_{i,1} \\ \mathbf{u}_{\Gamma_1,1} \\ \mathbf{u}_{\Gamma_2,1} \\ \hdashline \mathbf{q}_{i,2} \\ \mathbf{u}_{\Gamma_2,2} \\ \mathbf{u}_{\Gamma_3,2} \\ \hdashline \mathbf{q}_{i,3} \\ \mathbf{u}_{\Gamma_1,3} \\ \hdashline \mathbf{q}_{i,4} \\ \mathbf{u}_{\Gamma_2,4} \\ \hdashline \mathbf{q}_{i,5} \\ \mathbf{u}_{\Gamma_3,5} \end{Bmatrix} = \underbrace{\left[\begin{array}{cccccccc} \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \hdashline \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} \\ \hdashline \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \hdashline \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \hdashline \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} \end{array}\right]}_{\mathbf{L}^{\text{HCB}}} \begin{Bmatrix} \mathbf{q}_{i,1} \\ \mathbf{q}_{i,2} \\ \mathbf{q}_{i,3} \\ \mathbf{q}_{i,4} \\ \mathbf{q}_{i,5} \\ \mathbf{u}_{\Gamma_1} \\ \mathbf{u}_{\Gamma_2} \\ \mathbf{u}_{\Gamma_3} \end{Bmatrix}, \tag{9}$$

where the boundary DOF of each substructure are further partitioned according to the interface sets Γ_1, Γ_2, and Γ_3 shown in Figure 2. Finally, the assembled HCB mass and stiffness matrices are formed by adding the coupled contribution from each substructure as

$$\mathbf{M}^{\text{HCB}} = \sum_{j=1}^{n} \left(\mathbf{L}_j^{\text{HCB}}\right)^{\text{T}} \mathbf{M}_j^{\text{HCB}} \mathbf{L}_j^{\text{HCB}}, \qquad \mathbf{K}^{\text{HCB}} = \sum_{j=1}^{n} \left(\mathbf{L}_j^{\text{HCB}}\right)^{\text{T}} \mathbf{K}_j^{\text{HCB}} \mathbf{L}_j^{\text{HCB}}, \tag{10}$$

where the j^{th} substructure coupling matrix, $\mathbf{L}_j^{\text{HCB}}$, is formed by extracting the corresponding rows of $\mathbf{L}^{\text{HCB}}$. In this and all following cases, the same transformation is applied to both the mass and stiffness matrices; from this point forward only the transformation for the mass matrix will be shown for brevity.

2.2. System-Level Reduction

Craig and Chang [9], and later Castanier et al. [5], presented the characteristic constraint mode reduction as a system-level approach to reduce the size of the interface DOF of the assembled HCB model. This technique utilizes a secondary eigenvalue analysis on the interface partition of the assembled system

$$\left(\mathbf{K}_{bb}^{\text{HCB}} - \omega_r^2 \mathbf{M}_{bb}^{\text{HCB}}\right) \boldsymbol{\phi}_r^{\text{S-CC}} = \mathbf{0}. \tag{11}$$

The eigenvectors resulting from Eq. (11) are referred to as S-CC modes. A truncated set of these mass normalized S-CC modes are collected into a mode shape matrix,

$$\boldsymbol{\Phi}^{\text{S-CC}} = \left[\boldsymbol{\phi}_1^{\text{S-CC}}, \ldots, \boldsymbol{\phi}_n^{\text{S-CC}}\right]. \tag{12}$$

The truncated modal matrix $\boldsymbol{\Phi}^{\text{S-CC}}$ is used to transform the HCB interface DOF onto the S-CC generalized interface DOF,

$$\mathbf{u}_b \approx \boldsymbol{\Phi}^{\text{S-CC}}\mathbf{q}_b. \tag{13}$$

Pre- and post- multiplication of the assembled HCB matrices in Eq. (10) with an S-CC transformation matrix leads to the S-CC reduced mass and stiffness matrices

$$\mathbf{M}^{\text{S-CC}} = \underbrace{\begin{bmatrix}\mathbf{I} & \mathbf{0}\\ \mathbf{0} & \boldsymbol{\Phi}^{\text{S-CC}}\end{bmatrix}^{\text{T}}\begin{bmatrix}\mathbf{M}_{ii}^{\text{HCB}} & \mathbf{M}_{ib}^{\text{HCB}}\\ \mathbf{M}_{bi}^{\text{HCB}} & \mathbf{M}_{bb}^{\text{HCB}}\end{bmatrix}\begin{bmatrix}\mathbf{I} & \mathbf{0}\\ \mathbf{0} & \boldsymbol{\Phi}^{\text{S-CC}}\end{bmatrix}}_{\mathbf{T}^{\text{S-CC}}}. \tag{14}$$

The size of the system interface partition is reduced from the number of physical DOF on the interface to the number of S-CC modes in the truncated eigenvector set, $n_{\boldsymbol{\Phi}_b}$.

2.3. Local-Level Reduction

The system-level reduction introduced in Section 2.2 performs a secondary eigenvalue analysis on the interface partition of the assembled stiffness and mass matrices. In this section, four local-level interface reduction techniques are described, where all the interface reduction steps are taken before the assembly of subsystems.

2.3.1. Exact Compatibility [12]

The EC L-CC interface reduction proposed by Hong et al. [12] is based on a secondary eigenvalue analysis of the subcomponent HCB model's interface. Unlike the S-CC reduction, this is performed locally for each individual substructure rather than on the assembled matrices. For the j^{th} subcomponent, the secondary eigenvalue analysis on the interface DOF is written as

$$\left(\mathbf{K}_{bb,j}^{\text{HCB}} - \omega_r^2\mathbf{M}_{bb,j}^{\text{HCB}}\right)\left\{\boldsymbol{\phi}_j^{\text{L-CC}}\right\}_r = \mathbf{0}, \tag{15}$$

A truncated set of mass normalized eigenvector solutions collected in $\boldsymbol{\Phi}_j^{\text{L-CC}}$ denote the L-CC modes of the j^{th} substructure. The truncation is performed using frequency-based mode selection. The L-CC modes for each substructure are then used to reduce the interface DOF for each interface set locally.

The W-bracket system in Fig. 2 is used to illustrate the local interface reduction technique, as done by Hong et al. [12]. The L-CC modes in $\boldsymbol{\Phi}_1^{\text{L-CC}}$ from Eq. (15) have interface DOF for boundary sets $\boldsymbol{\Gamma}_1$ and $\boldsymbol{\Gamma}_2$, while the modes in $\boldsymbol{\Phi}_2^{\text{L-CC}}$ have interface DOF for boundary sets $\boldsymbol{\Gamma}_2$ and $\boldsymbol{\Gamma}_3$. Substructures 3, 4 and 5 only have one interface $\boldsymbol{\Gamma}_1$, $\boldsymbol{\Gamma}_2$ and $\boldsymbol{\Gamma}_3$, respectively. The L-CC modes for each subcomponent (1 through 5) are expressed then as

$$\begin{aligned}
&\boldsymbol{\Phi}_1^{\text{L-CC}} = \begin{bmatrix}\boldsymbol{\Phi}_{\Gamma_1,1}^{\text{L-CC}}\\ \boldsymbol{\Phi}_{\Gamma_2,1}^{\text{L-CC}}\end{bmatrix}, \quad \boldsymbol{\Phi}_2^{\text{L-CC}} = \begin{bmatrix}\boldsymbol{\Phi}_{\Gamma_2,2}^{\text{L-CC}}\\ \boldsymbol{\Phi}_{\Gamma_3,2}^{\text{L-CC}}\end{bmatrix},\\
&\boldsymbol{\Phi}_3^{\text{L-CC}} = \boldsymbol{\Phi}_{\Gamma_1,3}^{\text{L-CC}}, \quad \boldsymbol{\Phi}_4^{\text{L-CC}} = \boldsymbol{\Phi}_{\Gamma_2,4}^{\text{L-CC}}, \quad \boldsymbol{\Phi}_5^{\text{L-CC}} = \boldsymbol{\Phi}_{\Gamma_3,5}^{\text{L-CC}}.
\end{aligned} \tag{16}$$

The substructure interface modes are computed in Eq. (15) for the entire interface set simultaneously even if the interface is made up of multiple distinct interfaces. For each distinct interface, an augmented set of L-CC modes is established by collecting the modal components from every substructure along that interface. For the W-bracket system in Fig. 2, it holds that

$$\begin{aligned}
&\boldsymbol{\Phi}_{\boldsymbol{\Gamma}_1}^{\text{aug}} = \begin{bmatrix}\boldsymbol{\Phi}_{\Gamma_1,1}^{\text{L-CC}} & \boldsymbol{\Phi}_{\Gamma_1,3}^{\text{L-CC}}\end{bmatrix}, \quad \boldsymbol{\Phi}_{\boldsymbol{\Gamma}_2}^{\text{aug}} = \begin{bmatrix}\boldsymbol{\Phi}_{\Gamma_2,1}^{\text{L-CC}} & \boldsymbol{\Phi}_{\Gamma_2,2}^{\text{L-CC}} & \boldsymbol{\Phi}_{\Gamma_2,4}^{\text{L-CC}}\end{bmatrix},\\
&\boldsymbol{\Phi}_{\boldsymbol{\Gamma}_3}^{\text{aug}} = \begin{bmatrix}\boldsymbol{\Phi}_{\Gamma_3,2}^{\text{L-CC}} & \boldsymbol{\Phi}_{\Gamma_3,5}^{\text{L-CC}}\end{bmatrix}.
\end{aligned} \tag{17}$$

The augmented modes for a particular interface can represent the interface dynamics of the assembled system quite well, even though they are obtained from the eigenvectors of the uncoupled systems. The augmented interface L-CC

modes for each interface may lead to a set of vectors that are not necessarily linearly independent. To prevent ill conditioning, singular value decomposition (SVD) reformulates the L-CC modes and guarantees orthogonality,

$$\mathbf{U}^{\text{aug}}_{\Gamma_j}\mathbf{S}^{\text{aug}}_{\Gamma_j}\left(\mathbf{V}^{\text{aug}}_{\Gamma_j}\right)^{\text{T}} = \mathbf{\Phi}^{\text{aug}}_{\Gamma_j}, \tag{18}$$

where $\mathbf{U}^{\text{aug}}_{\Gamma_j}$ and $\mathbf{S}^{\text{aug}}_{\Gamma_j}$ are the left singular vectors and diagonal singular value matrix for the j^{th} interface set. The singular values describe how important each left singular vector is in describing the original basis. So if some singular values are very small, the corresponding left singular vectors can be truncated without significantly affecting the basis. A range-space truncation tolerance, σ_R, may be specified to determine the singular value cutoff as a fraction of the maximum singular value. In Hong's analysis [12], the truncation tolerance used was $\sigma_R = 1 \times 10^{-4} = 0.01\%$. This means that only the left singular vectors corresponding to singular values larger than 0.01% of the maximum singular value are kept for each interface set, and placed in $\mathbf{U}_{\mathbf{\Gamma}_j}$. Using an optimal value for σ_R will produce smaller model sizes without significantly affecting the accuracy of the solution. Setting σ_R too large will truncate too many modes and increase model error, while a small σ_R will result in more interface DOF than necessary to achieve practically the same accuracy. Appendix A.1 examines the effect of σ_R on the model performance.

The resorted matrices $\mathbf{U}_{\mathbf{\Gamma}_j}$ are then grouped for each subsystem j to form a block-diagonal matrix that contains all the interface sets for substructure j. Take substructures 1 and 3 as examples. The orthogonal L-CC basis is expressed as

$$\mathbf{U}_1 = \begin{bmatrix} \mathbf{U}_{\Gamma_1} & \mathbf{0} \\ \mathbf{0} & \mathbf{U}_{\Gamma_2} \end{bmatrix}, \quad \mathbf{U}_3 = \mathbf{U}_{\Gamma_1}. \tag{19}$$

Then, the reduced mass matrix is given for each subcomponent j as

$$\mathbf{M}^{\text{EC L-CC}}_j = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{U}_j \end{bmatrix}^{\text{T}} \begin{bmatrix} \mathbf{M}^{\text{HCB}}_{ii,j} & \mathbf{M}^{\text{HCB}}_{ib,j} \\ \mathbf{M}^{\text{HCB}}_{bi,j} & \mathbf{M}^{\text{HCB}}_{bb,j} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{U}_j \end{bmatrix}, \tag{20}$$

where the superscript EC L-CC indicates that the matrices are constructed using the local interface reduction technique with exact compatibility, as in [12].

Geometric compatibility conditions are enforced by relating the substructure interface DOF to a set of system interface DOF through a coupling matrix. The coupling matrix for the W-bracket system is written as,

$$\begin{Bmatrix} \mathbf{q}_{i,1} \\ \mathbf{q}_{\Gamma_1,1} \\ \mathbf{q}_{\Gamma_2,1} \\ \hdashline \mathbf{q}_{i,2} \\ \mathbf{q}_{\Gamma_2,2} \\ \mathbf{q}_{\Gamma_3,2} \\ \hdashline \mathbf{q}_{i,3} \\ \mathbf{q}_{\Gamma_1,3} \\ \hdashline \mathbf{q}_{i,4} \\ \mathbf{q}_{\Gamma_2,4} \\ \hdashline \mathbf{q}_{i,5} \\ \mathbf{q}_{\Gamma_3,5} \end{Bmatrix} = \underbrace{\left[\begin{array}{cccccccc} \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \hdashline \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} \\ \hdashline \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \hdashline \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \hdashline \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} \end{array}\right]}_{\mathbf{L}^{\text{EC L-CC}}} \begin{Bmatrix} \mathbf{q}_{i,1} \\ \mathbf{q}_{i,2} \\ \mathbf{q}_{i,3} \\ \mathbf{q}_{i,4} \\ \mathbf{q}_{i,5} \\ \mathbf{q}_{\Gamma_1} \\ \mathbf{q}_{\Gamma_2} \\ \mathbf{q}_{\Gamma_3} \end{Bmatrix}. \tag{21}$$

Finally, the reduced-order system is assembled by summing over the substructures as,

$$\mathbf{M}^{\text{EC L-CC}} = \sum_{j=1}^{n} \left(\mathbf{L}^{\text{EC L-CC}}_j\right)^{\text{T}} \mathbf{M}^{\text{EC L-CC}}_j \mathbf{L}^{\text{EC L-CC}}_j, \tag{22}$$

where $\mathbf{L}_j^{\text{EC L-CC}}$ denotes the rows of the coupling matrix corresponding to the j^{th} substructure.

2.3.2. Weak Compatibility

Similar to the EC L-CC method, the WC L-CC method begins by computing local interface modes for each substructure using Eq. (15). This transformation expresses the physical interface DOF as a linear combination of the interface modes,

$$\mathbf{u}_{b,j} \approx \mathbf{\Phi}_j^{\text{L-CC}} \mathbf{q}_{b,j}. \tag{23}$$

Then the reduced-order substructure mass matrix for each substructure is formed as

$$\mathbf{M}_j^{\text{WC L-CC}} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Phi}_j^{\text{L-CC}} \end{bmatrix}^{\text{T}} \begin{bmatrix} \mathbf{M}_{ii,j}^{\text{HCB}} & \mathbf{M}_{ib,j}^{\text{HCB}} \\ \mathbf{M}_{bi,j}^{\text{HCB}} & \mathbf{M}_{bb,j}^{\text{HCB}} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Phi}_j^{\text{L-CC}} \end{bmatrix}. \tag{24}$$

The resulting reduced substructures can no longer be coupled directly because, in general, the interface modes from one substructure will not match the interface modes from adjacent substructures. In fact, the interface mode sets of each substructure typically will not even span the same space. Traditional primal coupling of substructures is simple because each interface that must be joined has the exact same basis. This allows a direct assembly approach using coupling matrices, $\mathbf{L}$, that are trivial to derive, as seen in Eq. (9) and (21). When the primal coupling is not trivial, $\mathbf{L}$ can be obtained by first expressing constraint equations at the interface,

$$\mathbf{B}\mathbf{q}_{cat} = \mathbf{0} \tag{25}$$

where $\mathbf{q}_{cat}$ is a concatenated vector of generalized interface coordinates from the uncoupled substructures. The transformation from assembled interface coordinates, $\mathbf{q}_{cat} = \mathbf{L}_b \mathbf{q}_{asm}$, satisfies the constraint equations as long as $\mathbf{L}_b$ lies in the null space of $\mathbf{B}$. Note that the subscript b denotes that only the boundary partition of the coupling matrix is considered because the interior partition requires no coupling and thus consists simply of identity blocks. The general approach to finding the coupling matrix is to form $\mathbf{B}$ and compute its null space.

For the W-bracket example, the constraint equations to be satisfied are as follows,

$$\mathbf{u}_{\Gamma_1,1} = \mathbf{u}_{\Gamma_1,3}, \qquad \mathbf{u}_{\Gamma_2,1} = \mathbf{u}_{\Gamma_2,2}, \qquad \mathbf{u}_{\Gamma_2,1} = \mathbf{u}_{\Gamma_2,4}, \qquad \mathbf{u}_{\Gamma_3,2} = \mathbf{u}_{\Gamma_3,5}. \tag{26}$$

Collecting all of the constraint equations into a single matrix and expressing them in modal coordinates gives,

$$\underbrace{\begin{bmatrix} \mathbf{\Phi}_{\Gamma_1,1}^{\text{L-CC}} & \mathbf{0} & -\mathbf{\Phi}_{\Gamma_1,3}^{\text{L-CC}} & \mathbf{0} & \mathbf{0} \\ \mathbf{\Phi}_{\Gamma_2,1}^{\text{L-CC}} & -\mathbf{\Phi}_{\Gamma_2,2}^{\text{L-CC}} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{\Phi}_{\Gamma_2,1}^{\text{L-CC}} & \mathbf{0} & \mathbf{0} & -\mathbf{\Phi}_{\Gamma_2,4}^{\text{L-CC}} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Phi}_{\Gamma_3,2}^{\text{L-CC}} & \mathbf{0} & \mathbf{0} & -\mathbf{\Phi}_{\Gamma_3,5}^{\text{L-CC}} \end{bmatrix}}_{\mathbf{B}} \begin{Bmatrix} \mathbf{q}_{b,1} \\ \mathbf{q}_{b,2} \\ \mathbf{q}_{b,3} \\ \mathbf{q}_{b,4} \\ \mathbf{q}_{b,5} \end{Bmatrix} = \mathbf{0}. \tag{27}$$

If the bases at interfaces match, then $\mathbf{B}$ will have a clear null space, however in the WC L-CC method this typically is not the case. It is still possible to form an approximate null space using the SVD, $\mathbf{B} = \mathbf{U}_\mathbf{B}\mathbf{S}_\mathbf{B}\mathbf{V}_\mathbf{B}^{\text{T}}$. The approximate null space, $\mathbf{N}$, is computed by collecting the columns of $\mathbf{V}_\mathbf{B}$ whose singular values are less than a null space tolerance, σ_N. Each column of $\mathbf{N}$ can be thought of as a linear combination of the original interface modes that approximately satisfies compatibility, or a "compatible vector." As long as the singular values associated with each compatible vector are small, the compatibility will be good (though typically not exact). Each compatible vector will be governed by a generalized coordinate in the coupled system, while the portion of the local interface modes in the subspace not spanned by the compatible vectors is discarded. A high value of σ_N gives a large number of vectors that are deemed compatible and hence a large number of generalized coordinates to represent the interface dynamics very well, but the quality of the interface compatibility suffers because some of these vectors may be associated with relatively large singular values. Conversely, a low value of σ_N gives good interface compatibility but may reduce the number of generalized coordinates at the interface such that the dynamics can no longer be captured. Appendix A.2 investigates

the effect of the approximate null space tolerance on the performance of the reduction. Once an approximate null space has been obtained, it is placed into the boundary partitions of the overall coupling matrix. For the W-bracket model, this gives:

$$\begin{Bmatrix} \mathbf{q}_{i,1} \\ \mathbf{q}_{b,1} \\ \hdashline \mathbf{q}_{i,2} \\ \mathbf{q}_{b,2} \\ \hdashline \mathbf{q}_{i,3} \\ \mathbf{q}_{b,3} \\ \hdashline \mathbf{q}_{i,4} \\ \mathbf{q}_{b,4} \\ \hdashline \mathbf{q}_{i,5} \\ \mathbf{q}_{b,5} \end{Bmatrix} = \underbrace{\left[\begin{array}{cccccc} \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{N}_1 \\ \hdashline \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{N}_2 \\ \hdashline \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{N}_3 \\ \hdashline \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{N}_4 \\ \hdashline \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{N}_5 \end{array}\right]}_{\mathbf{L}^{\text{WC L-CC}}} \begin{Bmatrix} \mathbf{q}_{i,1} \\ \mathbf{q}_{i,2} \\ \mathbf{q}_{i,3} \\ \mathbf{q}_{i,4} \\ \mathbf{q}_{i,5} \\ \mathbf{q}_{b} \end{Bmatrix}. \tag{28}$$

where $\mathbf{N}_j$ denotes the rows of $\mathbf{N}$ corresponding to the interface DOF of the j^{th} substructure. Finally, the reduced-order system is assembled by summing over the substructures.

$$\mathbf{M}^{\text{WC L-CC}} = \sum_{j=1}^{n} \left(\mathbf{L}_j^{\text{WC L-CC}}\right)^{\mathrm{T}} \mathbf{M}_j^{\text{WC L-CC}} \mathbf{L}_j^{\text{WC L-CC}}, \tag{29}$$

where $\mathbf{L}_j^{\text{WC L-CC}}$ contains the rows of $\mathbf{L}^{\text{WC L-CC}}$ corresponding to the j^{th} uncoupled substructure.

2.3.3. Uncoupled Weak Compatibility

The weak compatibility technique described in the previous section results in a single coupled set of interface DOF for the system. This means that the different interface sets cannot move independently. Another approach to weak compatibility interface reduction is to decouple the interface modes before enforcing weak compatibility so that the interface sets *can* move independently. This increases the number of interface DOF in the system but can make the interface reduction easier to apply in some cases. The local interface modes are split up by interface set as shown in Eq. (16). By doing so, the interface modes may lose their linear independence so a QR orthonormalization is performed on the modes before they are used further. The (now orthogonal) interface modes are decoupled by arranging the components corresponding to each interface set on the block diagonals,

$$\begin{aligned} &\mathbf{\Phi}_1^{\text{UWC L-CC}} = \begin{bmatrix} \mathbf{\Phi}_{\Gamma_1,1}^{\text{L-CC}} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Phi}_{\Gamma_2,1}^{\text{L-CC}} \end{bmatrix}; \quad \mathbf{\Phi}_2^{\text{UWC L-CC}} = \begin{bmatrix} \mathbf{\Phi}_{\Gamma_2,2}^{\text{L-CC}} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Phi}_{\Gamma_3,2}^{\text{L-CC}} \end{bmatrix}, \\ &\mathbf{\Phi}_3^{\text{UWC L-CC}} = \mathbf{\Phi}_{\Gamma_1,3}^{\text{L-CC}}, \quad \mathbf{\Phi}_4^{\text{UWC L-CC}} = \mathbf{\Phi}_{\Gamma_2,4}^{\text{L-CC}}, \quad \mathbf{\Phi}_5^{\text{UWC L-CC}} = \mathbf{\Phi}_{\Gamma_3,5}^{\text{L-CC}}, \end{aligned} \tag{30}$$

Then, the reduced-order substructure mass matrix can be formed as they were in the previous sections,

$$\mathbf{M}_j^{\text{UWC L-CC}} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Phi}_j^{\text{UWC L-CC}} \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} \mathbf{M}_{ii,j}^{\text{HCB}} & \mathbf{M}_{ib,j}^{\text{HCB}} \\ \mathbf{M}_{bi,j}^{\text{HCB}} & \mathbf{M}_{bb,j}^{\text{HCB}} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{\Phi}_j^{\text{UWC L-CC}} \end{bmatrix}. \tag{31}$$

Weak compatibility must now be enforced separately at every interface set. For the W-bracket model, the constraint equations for the three interfaces are as follows,

$$\underbrace{\begin{bmatrix} \mathbf{\Phi}_{\Gamma_1,1}^{\text{L-CC}} & -\mathbf{\Phi}_{\Gamma_1,3}^{\text{L-CC}} \end{bmatrix}}_{\mathbf{B}_{\Gamma_1}} \begin{Bmatrix} \mathbf{q}_{b,1} \\ \mathbf{q}_{b,3} \end{Bmatrix} = \mathbf{0}, \tag{32}$$

$$\underbrace{\begin{bmatrix} \boldsymbol{\Phi}^{\text{L-CC}}_{\Gamma_2,1} & -\boldsymbol{\Phi}^{\text{L-CC}}_{\Gamma_2,2} & \mathbf{0} \\ \boldsymbol{\Phi}^{\text{L-CC}}_{\Gamma_2,1} & \mathbf{0} & -\boldsymbol{\Phi}^{\text{L-CC}}_{\Gamma_2,4} \end{bmatrix}}_{\mathbf{B}_{\Gamma_2}} \begin{Bmatrix} \mathbf{q}_{b,1} \\ \mathbf{q}_{b,2} \\ \mathbf{q}_{b,4} \end{Bmatrix} = \mathbf{0}, \tag{33}$$

$$\underbrace{\begin{bmatrix} \boldsymbol{\Phi}^{\text{L-CC}}_{\Gamma_3,2} & -\boldsymbol{\Phi}^{\text{L-CC}}_{\Gamma_3,5} \end{bmatrix}}_{\mathbf{B}_{\Gamma_3}} \begin{Bmatrix} \mathbf{q}_{b,2} \\ \mathbf{q}_{b,5} \end{Bmatrix} = \mathbf{0}. \tag{34}$$

Compatibility is enforced by computing the approximate null space of each compatibility matrix. Once again this requires an approximate null space tolerance, σ_N, to be defined. Appendix A.3 investigates the effect of varying this tolerance on the performance of the UWC L-CC method for the W-bracket model. The approximate null spaces from every interface are placed into the coupling transformation matrix,

$$\begin{Bmatrix} \mathbf{q}_{i,1} \\ \mathbf{q}_{\Gamma_1,1} \\ \mathbf{q}_{\Gamma_2,1} \\ \hdashline \mathbf{q}_{i,2} \\ \mathbf{q}_{\Gamma_2,2} \\ \mathbf{q}_{\Gamma_3,2} \\ \hdashline \mathbf{q}_{i,3} \\ \mathbf{q}_{\Gamma_1,3} \\ \hdashline \mathbf{q}_{i,4} \\ \mathbf{q}_{\Gamma_2,4} \\ \hdashline \mathbf{q}_{i,5} \\ \mathbf{q}_{\Gamma_3,5} \end{Bmatrix} = \underbrace{\left[\begin{array}{cccccccc} \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{N}_{\Gamma_1,1} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{N}_{\Gamma_2,1} & \mathbf{0} \\ \hdashline \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{N}_{\Gamma_2,2} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{N}_{\Gamma_3,2} \\ \hdashline \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{N}_{\Gamma_1,3} & \mathbf{0} & \mathbf{0} \\ \hdashline \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{N}_{\Gamma_2,4} & \mathbf{0} \\ \hdashline \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{N}_{\Gamma_3,5} \end{array}\right]}_{\mathbf{L}^{\text{UWC L-CC}}} \begin{Bmatrix} \mathbf{q}_{i,1} \\ \mathbf{q}_{i,2} \\ \mathbf{q}_{i,3} \\ \mathbf{q}_{i,4} \\ \mathbf{q}_{i,5} \\ \mathbf{q}_{\Gamma_1} \\ \mathbf{q}_{\Gamma_2} \\ \mathbf{q}_{\Gamma_3} \end{Bmatrix}. \tag{35}$$

where $\mathbf{N}_{\Gamma_j,i}$ denotes the rows of j^{th} interface's approximate null space corresponding to the ith substructure's interface DOF. The reduced-order system is assembled by summing over the substructures,

$$\mathbf{M}^{\text{UWC L-CC}} = \sum_{j=1}^{n} \left(\mathbf{L}_j^{\text{UWC L-CC}}\right)^{\text{T}} \mathbf{M}_j^{\text{UWC L-CC}} \mathbf{L}_j^{\text{UWC L-CC}}. \tag{36}$$

2.3.4. Virtual Node

Another local-level reduction technique was proposed by Lindberg et al. [19] where the interface DOF at the substructure are reduced with a set of undeformed coupling interface modes. This method was proposed as a way to connect stiff and soft substructures at a single VN having six DOF. This subsection reviews the interface reduction technique and generalizes the approach to allow n virtual nodes along an interface.

When a group of nodes along an interface are tied to a VN with multi-point constraints, the node set can collectively move in six shapes defined by three orthogonal translations and three rotations around the position of the VN. To define the transformation from physical DOF to the VN DOF, first consider the transformation to the l^{th} interface node from the VN to which it is being tied as,

$$\begin{Bmatrix} u_l \\ v_l \\ w_l \\ \theta_{x,l} \\ \theta_{y,l} \\ \theta_{z,l} \end{Bmatrix} = \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 & \Delta Z & -\Delta Y \\ 0 & 1 & 0 & -\Delta Z & 0 & \Delta X \\ 0 & 0 & 1 & \Delta Y & -\Delta X & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}}_{\boldsymbol{\theta}_{\Gamma,l}} \begin{Bmatrix} u^{\mathrm{VN}} \\ v^{\mathrm{VN}} \\ w^{\mathrm{VN}} \\ \theta_x^{\mathrm{VN}} \\ \theta_y^{\mathrm{VN}} \\ \theta_z^{\mathrm{VN}} \end{Bmatrix}, \tag{37}$$

where

$$\Delta X = x_l - x_{\mathrm{VN}}, \quad \Delta Y = y_l - y_{\mathrm{VN}}, \quad \Delta Z = z_l - z_{\mathrm{VN}}, \tag{38}$$

l is an index over all physical boundary nodes being tied to the current VN, and x, y, z are the nodal coordinates. Note that if the original set of interface DOF do not feature some rotational or translational DOF (as for instance in solid elements), then the corresponding rows of $\boldsymbol{\theta}_{\Gamma,l}$ can simply be removed.

A transformation for the entire set of interface DOF is generated by concatenating the transformations formed for each node as

$$\boldsymbol{\theta}_{\Gamma} = \begin{bmatrix} \boldsymbol{\theta}_{\Gamma,1} \\ \vdots \\ \boldsymbol{\theta}_{\Gamma,n_b} \end{bmatrix}, \tag{39}$$

where n_b is the number of nodes that are being tied to the current virtual node.

An analogous transformation is formed for every VN defined along an interface. Then, an overall transformation is formed for each substructure by collecting the transformations from each virtual node in a block diagonal matrix. This is most clearly demonstrated using the W-bracket example. Assume that for this example, each interface set is represented with a single VN, for a total of three VNs in the system as

$$\boldsymbol{\Theta}_1 = \begin{bmatrix} \theta_{\Gamma_1} & \mathbf{0} \\ \mathbf{0} & \theta_{\Gamma_2} \end{bmatrix}, \quad \boldsymbol{\Theta}_2 = \begin{bmatrix} \theta_{\Gamma_2} & \mathbf{0} \\ \mathbf{0} & \theta_{\Gamma_3} \end{bmatrix}, \quad \boldsymbol{\Theta}_3 = \theta_{\Gamma_1}, \quad \boldsymbol{\Theta}_4 = \theta_{\Gamma_2}, \quad \boldsymbol{\Theta}_5 = \theta_{\Gamma_3}. \tag{40}$$

Then, the reduced-order substructure mass matrix can be formed once again as they were in the previous sections,

$$\mathbf{M}_j^{\mathrm{VN}} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Theta}_j \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} \mathbf{M}_{ii,j}^{\mathrm{HCB}} & \mathbf{M}_{ib,j}^{\mathrm{HCB}} \\ \mathbf{M}_{bi,j}^{\mathrm{HCB}} & \mathbf{M}_{bb,j}^{\mathrm{HCB}} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Theta}_j \end{bmatrix}. \tag{41}$$

The coupling matrix for the W-bracket system is as follows,

$$\begin{Bmatrix} \mathbf{q}_{i,1} \\ \mathbf{u}_{\Gamma_1,1}^{\mathrm{VN}} \\ \mathbf{u}_{\Gamma_2,1}^{\mathrm{VN}} \\ \hdashline \mathbf{q}_{i,2} \\ \mathbf{u}_{\Gamma_2,2}^{\mathrm{VN}} \\ \mathbf{u}_{\Gamma_3,2}^{\mathrm{VN}} \\ \hdashline \mathbf{q}_{i,3} \\ \mathbf{u}_{\Gamma_1,3}^{\mathrm{VN}} \\ \hdashline \mathbf{q}_{i,4} \\ \mathbf{u}_{\Gamma_2,4}^{\mathrm{VN}} \\ \hdashline \mathbf{q}_{i,5} \\ \mathbf{u}_{\Gamma_3,5}^{\mathrm{VN}} \end{Bmatrix} = \underbrace{\left[\begin{array}{cccccccc} \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \hdashline \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} \\ \hdashline \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \hdashline \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \hdashline \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} \end{array}\right]}_{\mathbf{L}^{\mathrm{VN}}} \begin{Bmatrix} \mathbf{q}_{i,1} \\ \mathbf{q}_{i,2} \\ \mathbf{q}_{i,3} \\ \mathbf{q}_{i,4} \\ \mathbf{q}_{i,5} \\ \mathbf{u}_{\Gamma_1}^{\mathrm{VN}} \\ \mathbf{u}_{\Gamma_2}^{\mathrm{VN}} \\ \mathbf{u}_{\Gamma_3}^{\mathrm{VN}} \end{Bmatrix}. \tag{42}$$

The j^{th} substructure coupling matrix, $\mathbf{L}_j^{\mathrm{VN}}$, is formed by extracting the corresponding rows of $\mathbf{L}^{\mathrm{VN}}$. Finally, the coupled VN mass matrix is formed by adding the contribution from each substructure,

$$\mathbf{M}^{\mathrm{VN}} = \sum_{j=1}^{n} \left(\mathbf{L}_j^{\mathrm{VN}}\right)^{\mathrm{T}} \mathbf{M}_j^{\mathrm{VN}} \mathbf{L}_j^{\mathrm{VN}}. \tag{43}$$

The selection of VN positions can be performed manually but in this work it is automated using the k-means node clustering algorithm to produce groups of nodes that are closely spaced. This enabled the number of VNs, k, to be varied easily so that accuracy with respect to k could be investigated. There is no general guideline for selecting k, but for certain problems, physical intuition may be used to determine groups of interface nodes that will behave nearly rigidly, and can thus be represented by VNs. In the can-beam problem shown in section 3.3 for example, each bolted connection is represented with a single VN. If it is not possible to manually identify rigid node groups, then the VN method is likely not the best choice of interface reduction anyway.

Once interface nodes are grouped, the VN for each group is computed by averaging the nodal coordinates for all nodes in that group. The k-means algorithm uses an optimization strategy that can get stuck in local minima, so several restarts may be necessary. In this work, 100 restarts are used in selecting the groups and the k-means algorithm is applied to the nodes from a single interface at a time to ensure that one VN will not be tied to nodes from different interfaces.

2.4. Hybrid Reduction

As shown later in Section 3, the S-CC interface reduction in Section 2.2 produces the most accurate results since the secondary eigenvalue analysis accounts for all the mass and stiffness of the interface DOF. Unfortunately, the advantage of local modal analysis is lost with this substructuring technique. The local-level interface reductions presented in Section 2.3 retain this advantage, however the accuracy of the global reduced-order model suffers and requires more modes to match the system-level approach. This is due to the fact that the mass and stiffness coupling between different substructures is totally neglected during the local modal analysis and must be enforced approximately.

A novel hybrid interface reduction technique is proposed here where the secondary eigenvalue analysis is performed on a single interface shared by a set of substructures. By doing so, the coupling between neighboring substructures, which is of great significance, can be taken into consideration, but the size of the eigenvalue problem is reduced to only the size of each individual interface. In the limiting case of two substructures being assembled at one common interface (e.g. the Cantilever-Plate Model in Section 3.1), this method is exactly the same as the S-CC interface reduction.

The W-bracket example shows how the hybrid-level interface reduction is applied at an individual interface. The secondary modal analysis for all the substructures connecting to the j^{th} interface Γ_j is written as

$$\left(\mathbf{K}_{bb}^{\mathrm{HCB},\Gamma_j} - \omega_r^2 \mathbf{M}_{bb}^{\mathrm{HCB},\Gamma_j}\right) \boldsymbol{\phi}_r^{\mathrm{H\text{-}CC},\Gamma_j} = \mathbf{0} \tag{44}$$

where the mass matrix is assembled only for the substructures connecting to interface Γ_j,

$$\mathbf{M}^{\mathrm{HCB},\Gamma_j} = \sum_{k\in\Gamma_j} \left(\mathbf{L}_k^{\mathrm{HCB}}\right)^{\mathrm{T}} \mathbf{M}_k^{\mathrm{HCB}} \mathbf{L}_k^{\mathrm{HCB}}, \tag{45}$$

Note that Eq. (44) specifies that the entire boundary partition of the partially assembled mass and stiffness matrices, $\mathbf{M}^{\mathrm{HCB},\Gamma_j}$ and $\mathbf{K}^{\mathrm{HCB},\Gamma_j}$, be used in computing the eigenvectors. Thus, the mass normalized matrix $\boldsymbol{\Phi}^{\mathrm{H\text{-}CC},\Gamma_j}$, referred to as the hybrid characteristic constraint modes, contains DOF for all interface sets that are attached to substructures adjacent to Γ_j. Including these "extra" interface sets allows the free-interface modes of the subassembly to be computed. This is particularly important when rigid-body interface motion is present in the system modes. The interface modes computed in Eq. (44) for Γ_j may contain terms corresponding to other interfaces as well but can simply be discarded. For the W-bracket model, the H-CC modes then become

$$\boldsymbol{\Phi}^{\text{H-CC},\Gamma_1} = \begin{bmatrix} \boldsymbol{\Phi}_{\Gamma_1}^{\text{H-CC},\Gamma_1} \\ \cancel{\boldsymbol{\Phi}_{\Gamma_2}^{\text{H-CC},\Gamma_1}} \end{bmatrix}, \quad \boldsymbol{\Phi}^{\text{H-CC},\Gamma_2} = \begin{bmatrix} \cancel{\boldsymbol{\Phi}_{\Gamma_1}^{\text{H-CC},\Gamma_2}} \\ \boldsymbol{\Phi}_{\Gamma_2}^{\text{H-CC},\Gamma_2} \\ \cancel{\boldsymbol{\Phi}_{\Gamma_3}^{\text{H-CC},\Gamma_2}} \end{bmatrix}, \quad \boldsymbol{\Phi}^{\text{H-CC},\Gamma_3} = \begin{bmatrix} \cancel{\boldsymbol{\Phi}_{\Gamma_2}^{\text{H-CC},\Gamma_3}} \\ \boldsymbol{\Phi}_{\Gamma_3}^{\text{H-CC},\Gamma_3} \end{bmatrix}. \tag{46}$$

Discarding part of a mode shape can cause the remaining component to lose its linear dependence, so the retained portions are orthonormalized with a QR decomposition to avoid ill-conditioning. The transformation from the HCB coordinates to the H-CC generalized coordinates is

$$\begin{Bmatrix} \mathbf{q}_{i,1} \\ \mathbf{q}_{i,2} \\ \mathbf{q}_{i,3} \\ \mathbf{q}_{i,4} \\ \mathbf{q}_{i,5} \\ \mathbf{u}_b^{\Gamma_1} \\ \mathbf{u}_b^{\Gamma_2} \\ \mathbf{u}_b^{\Gamma_3} \end{Bmatrix} = \underbrace{\begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{I} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \boldsymbol{\Phi}_{\Gamma_1}^{\text{H-CC},\Gamma_1} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \boldsymbol{\Phi}_{\Gamma_2}^{\text{H-CC},\Gamma_2} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \boldsymbol{\Phi}_{\Gamma_3}^{\text{H-CC},\Gamma_3} \end{bmatrix}}_{\mathbf{T}^{\text{H-CC}}} \begin{Bmatrix} \mathbf{q}_{i,1} \\ \mathbf{q}_{i,2} \\ \mathbf{q}_{i,3} \\ \mathbf{q}_{i,4} \\ \mathbf{q}_{i,5} \\ \mathbf{q}_{\Gamma_1} \\ \mathbf{q}_{\Gamma_2} \\ \mathbf{q}_{\Gamma_3} \end{Bmatrix}. \tag{47}$$

The H-CC reduced mass matrix reduces to

$$\mathbf{M}^{\text{H-CC}} = \left(\mathbf{T}^{\text{H-CC}}\right)^{\text{T}} \mathbf{M}^{\text{HCB}} \mathbf{T}^{\text{H-CC}}. \tag{48}$$

This hybrid interface reduction technique is most suitable for systems with numerous interface sets. When few substructures are added, removed or modified around a specified interface set, the secondary eigenvalue analysis and the coupling matrix need to be recalculated only for related substructures connecting to this interface. This is more efficient than the system-level approach and can reduce the computational cost without largely compromising accuracy.

3. Numerical Examples and Discussion

Three numerical examples are explored in the following subsections to evaluate the accuracy and efficiency of the interface reduction methods. Since all these methods operate on a nominal HCB substructured model, the number of fixed-interface modes for each case are chosen such that the maximum frequency error relative to the full fidelity model is less than 0.1% for the frequency range of interest. Furthermore, the SVD tolerance values used in all the examples are those recommended in the Appendix, namely $\sigma_R = 1 \times 10^{-4}$ for EC L-CC, and $\sigma_N = 1 \times 10^{-1}$ for WC and UWC L-CC.

Timing results are also included for each example. Note that the relative performance of the methods will depend on the model, the computing architecture, and the implementation of the reductions. The results are included to give the reader a sense of the computation times, but do not necessarily show the computational advantage of the reductions. The accuracy of the methods is investigated by using an eigenvalue solution to compute natural frequencies and mode shapes of each structure, but computationally this task is typically accomplished more quickly without resorting to model-order reduction. HCB reduced-order models are most effective at reducing large-scale, memory intensive problems and for accelerating different types of analyses, such as time simulations, nonlinear modal solutions, and cases where multiple configurations of substructures and boundary conditions must be considered. Showing timing results for these cases however is beyond the scope of this paper.

3.1. Cantilevered-Plate Model

The first numerical example under consideration is a cantilever-plate model consisting of two plates that share one interface as shown in Fig. 3a. The nodes along the long edge nearest to the viewer are fixed creating a cantilever boundary condition. The model features 1152 four-node shell elements, with three translational and three rotational DOF per node. After removing the DOF associated with the fixed nodes, the FE model has 7056 DOF. The dimensions and materials are the same as in [4, 5].

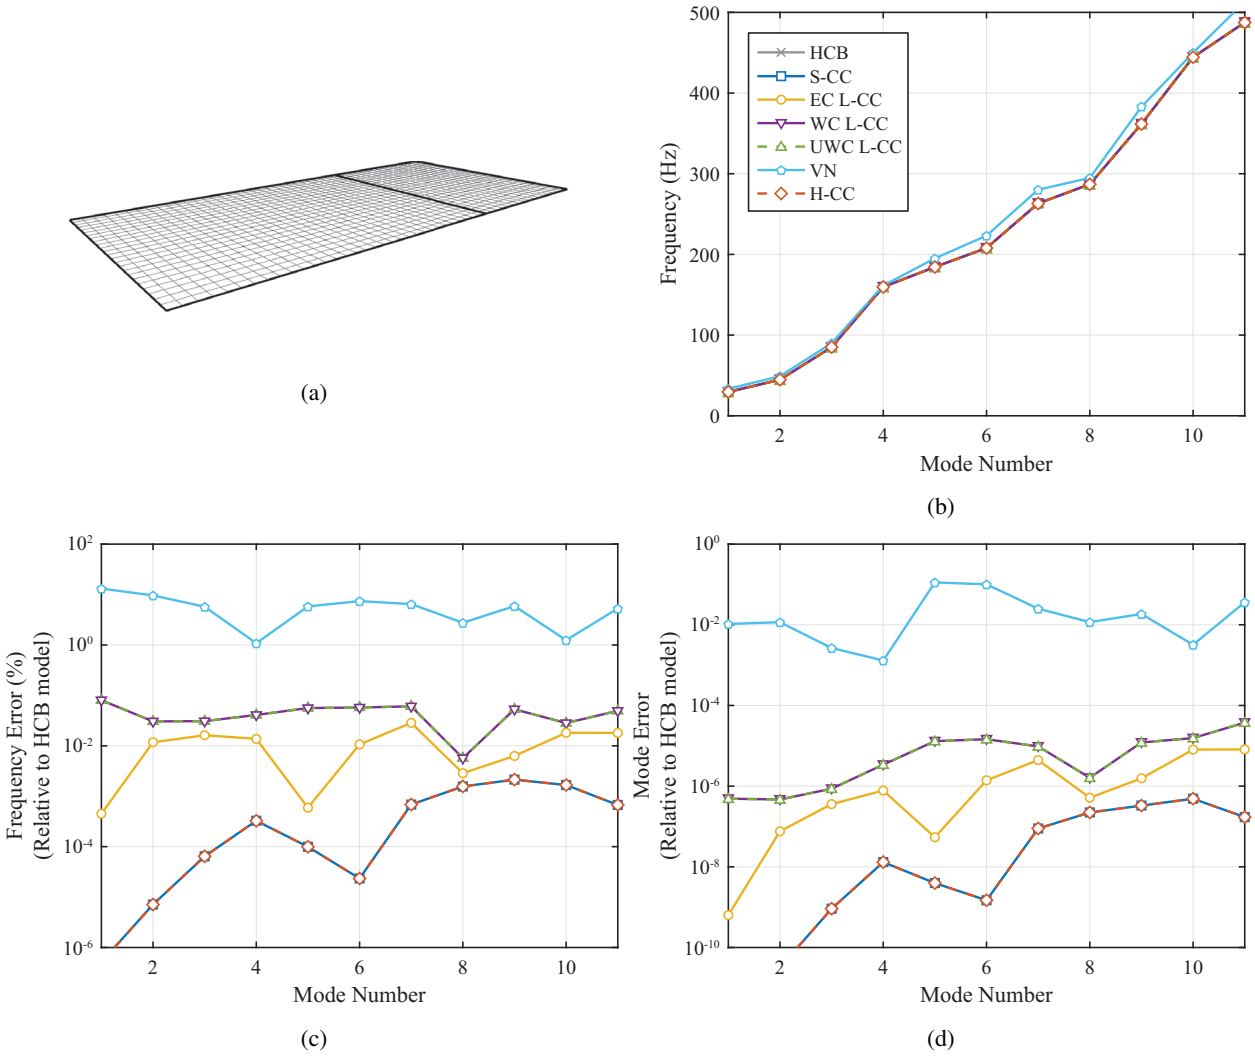

Figure 3: FE mesh (a), natural-frequency comparison (b), frequency-error comparison (c), and mode-error comparison (d) for the cantilever-plate model.

An HCB model of the cantilevered plate is created with a total of 27 fixed-interface modes, as summarized in Table 1, and 144 interface DOF. Each interface reduction is applied to the same base HCB model and then the modal solutions are computed from the assembled reduced-order models. The 11 natural frequencies below 500 Hz are plotted in Fig. 3b for all six interface reduction methods as well as for the reference HCB model. For this example, each interface-reduction method has a total of 18 interface DOF in its respective assembled model in order to make a fair comparison of error between the different methods. The number of interface modes and corresponding maximum frequency CC mode (where applicable) is shown in Table 2 for all the approaches under consideration. Note that because the cantilevered plate has just a single interface set, the H-CC method is identical to the S-CC method, and

the WC L-CC method becomes identical to the UWC L-CC method.

Table 1: The number of fixed-interface modes retained in each substructure in the cantilevered-plate model

Number of Internal Vibration Modes (Maximum Frequency, Hz)		
Sub. 1	Sub. 2	Total
18 (1440)	9 (1380)	27 (1440)

Table 2: The number (and maximum frequency where applicable) of interface modes used by each method in reducing the cantilevered-plate model

Method	Number of CC Modes Calculated per Substructure (Maximum Frequency, Hz)		Number of Interface DOF (Maximum Frequency, Hz)
	Sub. 1	Sub. 2	
HCB	-	-	144
S-CC / H-CC	-	-	18 (7600)
EC L-CC	9 (2540)	10 (3110)	18
WC L-CC / UWC L-CC	28 (19220)	28 (19030)	18
VN	-	-	3×6 = 18

Figures 3c and 3d show the error in natural frequencies and mode shapes, respectively, with respect to the HCB model. The mode error is defined as the deviation from unity of the diagonal components of the modal assurance criterion (MAC) matrix [1]. The two error plots reveal that some interface reduction methods provide better accuracy for the same (assembled) model size. The S-CC / H-CC methods show the lowest error, followed by the EC L-CC method, then the WC L-CC / UWC L-CC methods, and finally the VN method. Aside from VN, all methods comfortably achieve frequency errors below 0.1% for this example.

The computation times for the cantilevered plate model are summarized in Table 3. The model reduction time reported for each interface reduction is the time needed just to perform the interface reduction, and does not include the time needed to obtain the HCB model. The eigenvalue solution times are roughly equivalent for all of the interface reductions, which is expected because the final reduced-order model sizes are the same. Even though the reduced-order models obtained with the S-CC and H-CC methods are the same, the timing results are different because the codes have different amounts of overhead and there are also natural variations in timing data. The same observation can be made with the WC and UWC L-CC methods.

The timing results show that it is faster to simply perform the full model eigenvalue solution than it is to reduce the model and then compute a reduced-order eigenvalue solution. As mentioned earlier, this is expected because the eigenvalue analysis often does not expose the computational benefits of model-order reduction.

To gain more insight into the performance of each method (i.e., model size versus error), the frequency cut off value is incrementally increased to show the convergence of the maximum frequency error of the first 11 modes. The performance curves for each interface reduction method are shown in Fig. 4. These plots show how many interface DOF are required to obtain a given level of accuracy. For example, Table 4 gives the minimum number of interface

Table 3: Computation times needed to create each reduced-order model for the cantilever plate, and the computation time needed to solve for each reduce-order model's eigenvalues. (Note that the reduced-order modeling times for each interface reduction only include the time needed to compute the interface reduction and do not include the HCB modeling time.)

Method	Model Reduction Time (ms)	Eigenvalue Solution Time (ms)
Full FE	-	363
HCB	354	14
S-CC	18	8
H-CC	39	8
EC L-CC	71	9
WC L-CC	64	9
UWC L-CC	57	8
VN	109	8

DOF in the assembly and corresponding frequency cut off to get a maximum frequency error less than 0.1%. The frequency cut off is somewhat misleading when considering the EC L-CC, WC L-CC, and UWC L-CC methods because the modes are truncated at the substructure level and subsequently the modes are further operated on using an SVD. In most cases, the SVD discards some information from the mode set so that many or all of the modes below the frequency cut off are no longer exactly represented. Nevertheless, the underlying frequency cut off is given in Table 4 to provide readers with a sense of the differences between the methods.

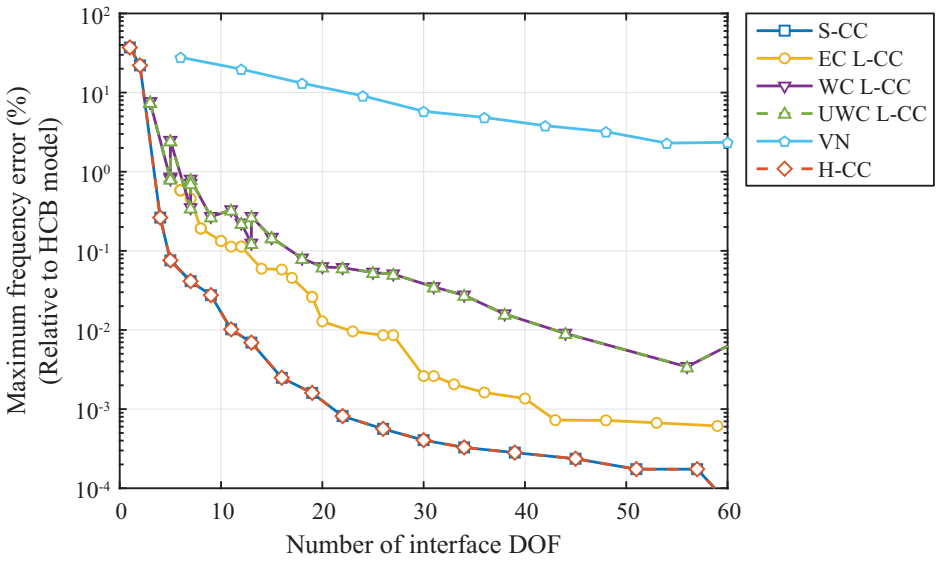

Figure 4: Study of maximum frequency error with increasing number of interface DOF for various interface reduction methods using the cantilevered-plate model

3.2. W-Bracket Model

The second numerical example considers the W-bracket model that was introduced previously. This model consists of five plates that are coupled together, as shown in Fig. 5a. The FE mesh consists of 2304 four-node shell elements with three displacements and three rotations per node, resulting in 14550 DOF for the whole model. The model has free-free boundary conditions allowing six rigid body modes.

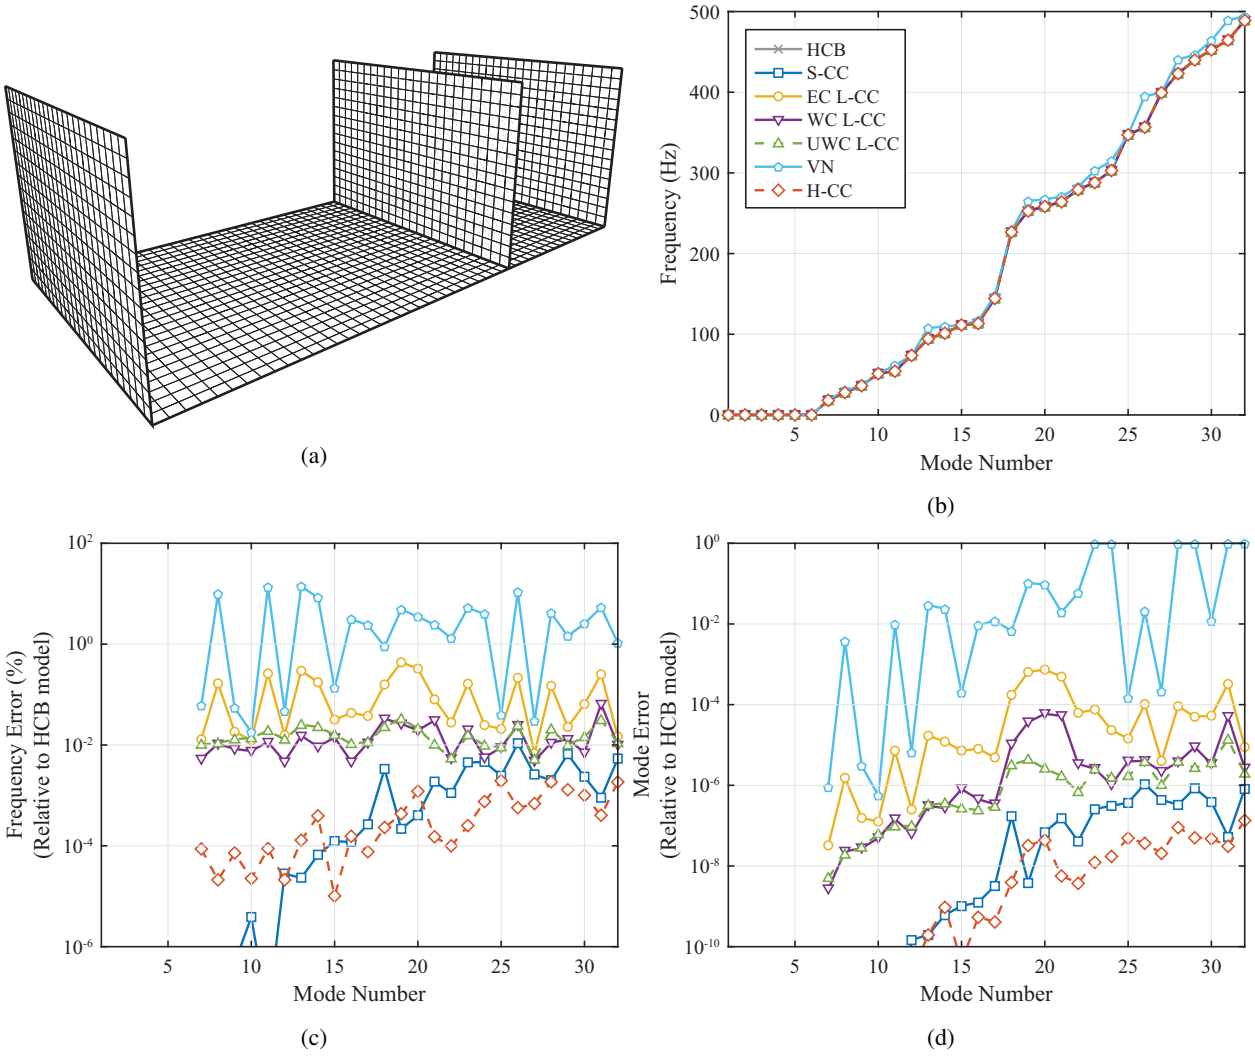

Figure 5: FE mesh (a), natural-frequency comparison (b), frequency-error comparison (c), and mode-error comparison (d) for the W-bracket model.

Table 4: Minimum interface size to achieve 0.1% frequency error for the cantilevered-plate model. *Note that the EC L-CC method uses 7 modes per substructure for a total of 14, and hence the maximum frequency of those L-CC modes is much lower than for the WC and UWC methods.

	S-CC / H-CC	EC L-CC	WC L-CC / UWC L-CC	VN
Number of CC Modes	5	14*	18	>60
Maximum Frequency, Hz	630	2030	19220	-

An HCB model of the W-bracket is formed with a total of 53 fixed-interface modes, as summarized in Table 5. Each of the three interfaces contains 150 DOF for a total of 450 interface DOF in the HCB model. The W-bracket has 32 vibration modes with frequencies below 500 Hz, including the six rigid-body modes. Figure 5b shows the computed natural frequencies of each interface reduction method along with the reference HCB model. For this example, each model has 54 interface DOF except for the EC L-CC method, which has 53. The number of DOF associated with each substructure and with each interface are shown in Table 6.

Table 5: The number of fixed-interface modes retained in the HCB model of each substructure in the W-bracket model

Number of Internal Vibration Modes (Maximum Frequency, Hz)					
Sub. 1	Sub. 2	Sub. 3	Sub. 4	Sub. 5	Total
16 (1201)	6 (1223)	10 (1172)	10 (1172)	11 (1241)	53 (1241)

Figures 5c and 5d respectively show the error in natural frequencies and mode shapes. Since the W-bracket has more than one distinct interface, the H-CC results differ from the S-CC results and the WC L-CC results differ from the UWC L-CC results. Once again, the S-CC and H-CC methods produce the lowest errors, with S-CC performing slightly better at lower frequencies but H-CC performing slightly better at higher frequencies. The next best performing methods are WC L-CC and UWC L-CC, both of which show relatively similar error levels. The EC L-CC method shows slightly higher error still, and the VN method again gives the highest error. Compared to the cantilevered-plate results, the main difference in performance is that the EC L-CC becomes less accurate relative to the WC and UWC L-CC techniques. This occurs because the number of interface DOF used by the EC L-CC method grows quickly when the number and complexity of interfaces is increased. Thus, in order to reach the same number of interface DOF as the other methods, the cut off frequency used for the EC L-CC method must be relatively low.

The computation times for the W-bracket model are summarized in Table 7. As with the cantilevered plate, the eigenvalue solution times are roughly equivalent for all of the interface reductions, which is expected because the final reduced-order model sizes are the same. Despite the more complex arrangement of substructures with multiple interfaces, the relative computation times mirror those from the cantilevered plate example.

As with the cantilevered-plate model, a performance plot is produced for the W-bracket model to compare error versus model size for each interface reduction. The maximum frequency error over the first 32 modes (excluding the rigid body modes) is evaluated with respect to the number of interface DOF for each reduced-order model as the cut off frequency is gradually incremented. The resulting curves are shown in Fig. 6 and mimics the performance seen in Figs. 5c and 5d. Table 8 shows the lowest number of interface DOF in Fig. 6 that achieves a maximum frequency error below 0.1%. Although the S-CC method requires the fewest interface DOF, the H-CC method has a much lower frequency cut-off. This suggests that comparing frequency cut-offs across different methods should be done with caution because the meaning depends on the method.

To understand the differences between the interface reduction methods, it helps to visualize the mode shapes. The first three elastic mode shapes of the assembled HCB model are plotted in Fig. 7 to give an idea of the types of interface motion that interface reduction modes must capture to achieve good accuracy. Note that although they are not shown here, the W-bracket model also admits rigid body modes which must also be recoverable by the interface reduction modes. Figure 8 shows three of the CM shapes used in the HCB model. The deflection in the substructure

Table 6: The number (and cutoff frequency where applicable) of interface modes used by each method in reducing the W-bracket model. *Note that for the EC, WC, and UWC methods the interface DOF are further truncated based on the singular value ratio; the number of modes and maximum frequencies per substructure do not reflect this.

Method	Number of CC modes Calculated per Substructure (Maximum Frequency, Hz)					Number of Interface DOF (Maximum Frequency, Hz)			
	Sub. 1	Sub. 2	Sub. 3	Sub. 4	Sub. 5	Γ_1	Γ_2	Γ_3	Total
HCB	-	-	-	-	-	150	150	150	450
S-CC	-	-	-	-	-	-	-	-	54 (11960)
EC L-CC*	15 (740)	13 (750)	9 (650)	9 (650)	9 (650)	16	21	16	53
WC L-CC*	64 (20930)	64 (21100)	32 (20850)	32 (20850)	32 (20850)	-	-	-	54
UWC L-CC*	58 (18090)	58 (18190)	29 (17980)	29 (17980)	29 (17980)	18	18	18	54
VN	-	-	-	-	-	3×6 = 18	3×6 = 18	3×6 = 18	54
H-CC	-	-	-	-	-	16 (1460)	23 (1460)	15 (1220)	54 (1460)

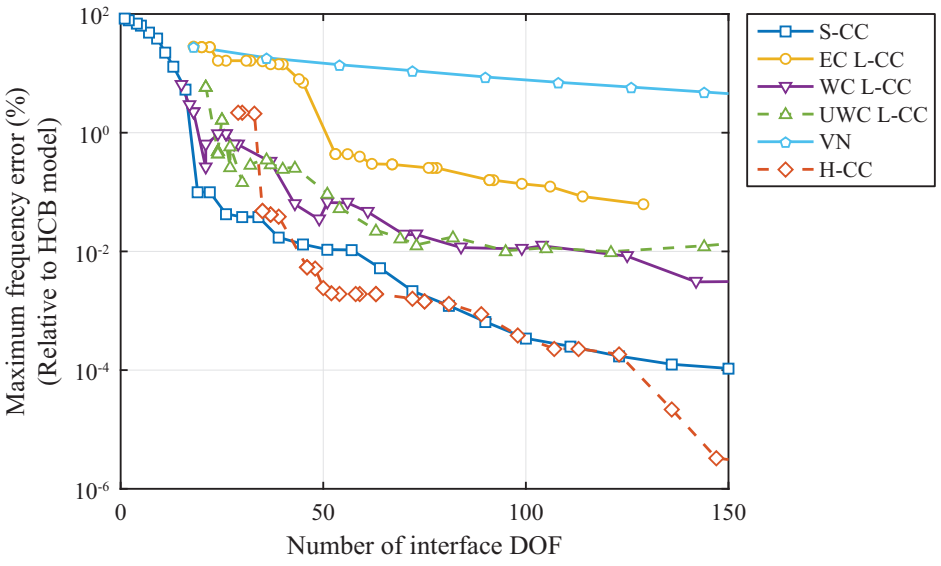

Figure 6: Study of maximum frequency error with increasing number of interface DOF for various interface reduction methods using the W-bracket model

Table 7: Computation times needed to create each reduced-order model for the W-bracket, and the computation time needed to solve for each reduce-order model's eigenvalues. (Note that the reduced-order modeling times for each interface reduction only include the time needed to compute the interface reduction and do not include the HCB modeling time.)

Method	Model Reduction Time (ms)	Eigenvalue Solution Time (ms)
Full FE	-	1026
HCB	965	98
S-CC	140	17
H-CC	221	15
EC L-CC	205	17
WC L-CC	448	15
UWC L-CC	409	18
VN	303	14

Table 8: Minimum interface size to achieve 0.1% frequency error for the W-bracket model

	S-CC	EC L-CC	WC L-CC	UWC L-CC	VN	H-CC
Number of CC Modes	26	114	43	51	>150	35
Maximum Frequency, Hz	3440	6200	14890	15360	-	280

interiors shows the static response to the CM's unit interface deflection while holding the rest of the interface DOF fixed. No individual CM captures the global behavior of the interface in Fig. 7 very well, but the linear combination of all of the CMs does. This type of interface representation is completely uncoupled, meaning that deflection of any interface DOF can occur independently of all other interface DOF.

Figure 9 shows select interface mode shapes for the S-CC, EC L-CC, WC L-CC, UWC L-CC, VN, and H-CC methods. All of these interface reduction mode shapes represent linear combinations of the CMs, such as the ones shown in Fig. 8. Thus, when these modes are projected back onto physical space, they also include deflection of the interior parts of every substructure. All the mode shapes except for the VN interface modes show smooth deflections compared to the CMs in Fig. 8 because they are all created using the characteristic vibration shapes of the interfaces in one way or another. In contrast, the VN mode shapes are simply rigid couplings of interface nodes which explains its poor performance on the cantilevered plate and W-bracket examples.

Among the six methods, the only two that show global coupling are the S-CC and the WC L-CC methods (Figs. 9a and 9c). This means that deflection of any interface DOF results in deformation throughout the entire model. The H-CC, EC L-CC, and UWC L-CC mode shapes have uncoupled interfaces (i.e., each interface mode involves deflection of a single interface and adjacent substructure interiors). The VN modes shows even less coupling, as small groups of nodes can be deflected independently. Generally, less coupling can provide increased versatility. For example, if interfaces are uncoupled it is easier to exchange substructures without updating the entire model, and the modal representation of a single interface can be refined. For the same final model size, uncoupled interface methods tend to produce smaller substructure model sizes. In terms of accuracy, uncoupled methods do not seem to have any disadvantage.

In this example, each of the methods is able to reproduce rigid body modes in the system and the S-CC and WC L-CC interface mode sets actually contain the global rigid body modes. For the other methods, the system rigid body mode shapes are represented using a linear combination of the interface modes. Note that the EC L-CC method uses a truncated SVD orthogonalization that could in theory discard part of the basis containing rigid body information. This was not observed in any of the example systems because the SVD truncation typically discards higher frequency modal information, but users should be aware of this possibility. If unexpected results occur with the EC L-CC method, then the SVD truncation can be replaced with a pure orthogonalization without any truncation.

The elastic system modes in general cannot be exactly recovered by any of the interface reduction methods,

however some methods give much better accuracy than others. The VN method will not achieve good accuracy when models have flexible interfaces like the W-bracket model. For the other five methods, it is not readily apparent from the mode shape diagrams why some outperform others. In general, characteristic shapes that are obtained from assembled interfaces are able to capture the dynamics of the assembled system better than characteristic shapes that are completely local.

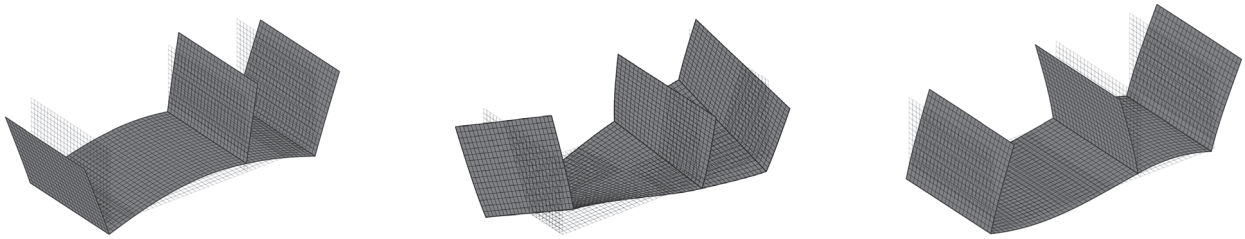

(a) Mode 7 (b) Mode 8 (c) Mode 9

Figure 7: First three elastic system mode shapes for the W-bracket, obtained with the HCB reduced-order model.

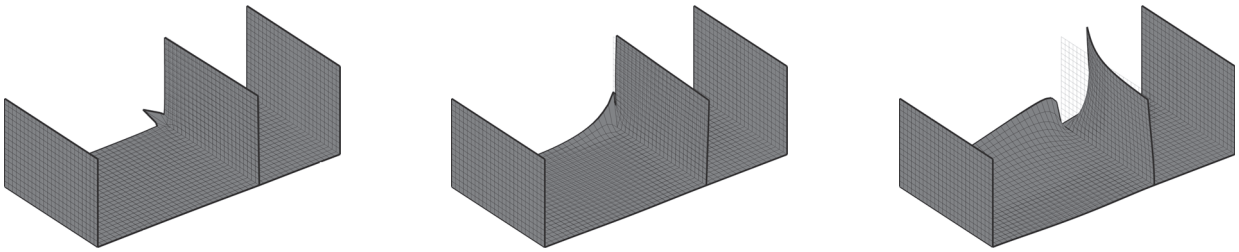

(a) Constraint Mode 2 (b) Constraint Mode 3 (c) Constraint Mode 5

Figure 8: Selected constraint modes from HCB representation.

3.3. Can-Beam Model

The third and final numerical example is performed with the can-beam model shown in Fig. 10a in order to demonstrate the methodologies on industrial sized models. This structure is split into two substructures: a can, and a beam attached to a plate, which are joined at the eight circular patches (shown in detail in the inset of Fig. 10a) to represent a bolted connection. The assembled FE model contains a mixture of 8-node shell elements and 20-node solid elements elements, resulting in a total of 137,955 DOF.

An HCB representation of system is created with 26 fixed-interface modes as summarized in Table 9. The can-beam model has 22 vibration modes with natural frequencies below 2000 Hz, including six rigid body modes. The natural frequencies are computed with the HCB model as well as with all six interface reduction models, and are plotted in Fig. 10b. Table 10 gives a summary of the number of interface DOF used by each method. For this example, each interface reduced model contains 48 interface DOF.

Table 9: The number of fixed-interface modes retained in the HCB model of each substructure in the can-beam model

Number of Internal Vibration Modes (Maximum Frequency, Hz)		
Sub. 1	Sub. 2	Total
20 (2360)	6 (1840)	26 (2360)

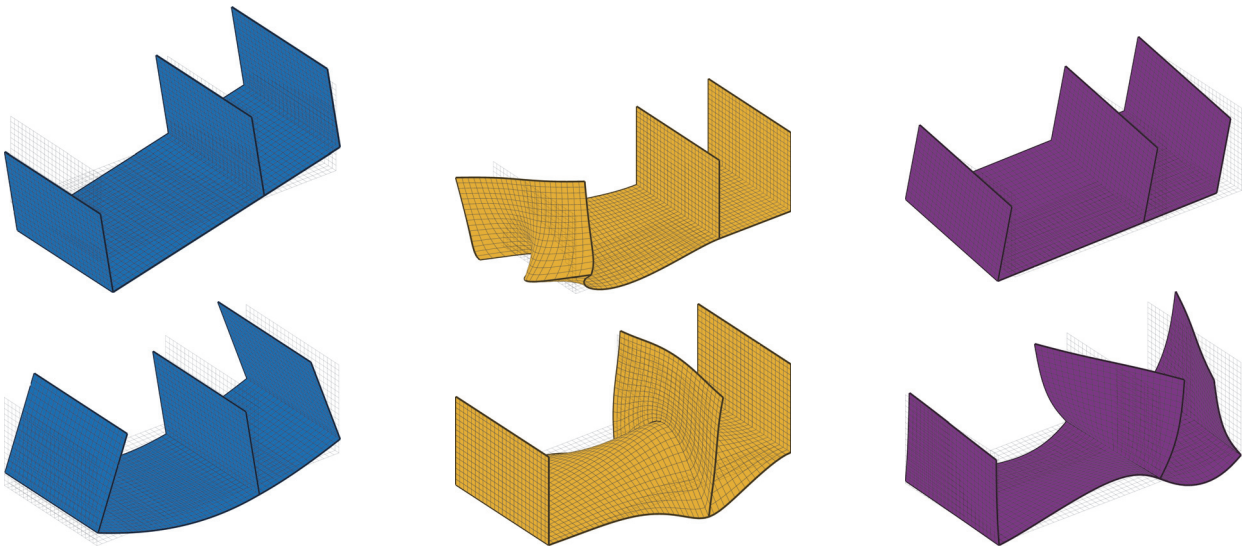

(a) S-CC (b) EC L-CC (c) WC L-CC

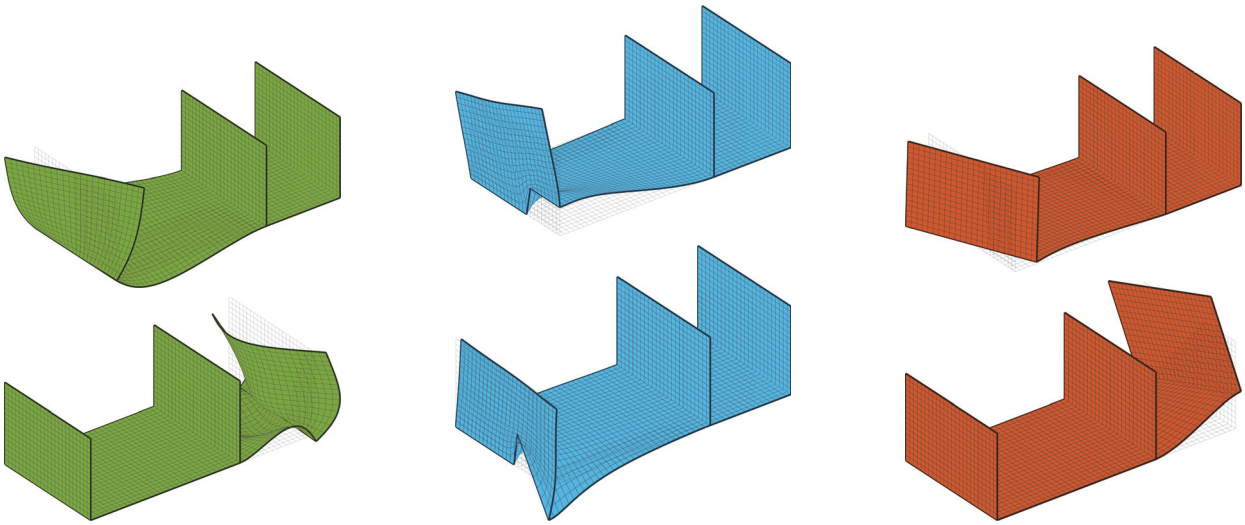

(d) UWC L-CC (e) VN (f) H-CC

Figure 9: Interface mode shapes for each interface reduction method. The S-CC and WC L-CC modes exhibit global deflection at all interface DOF. The EC L-CC, UWC L-CC, and H-CC modes show localized deformation at specific interface sets. The VN modes show localized deformation of certain node groups.

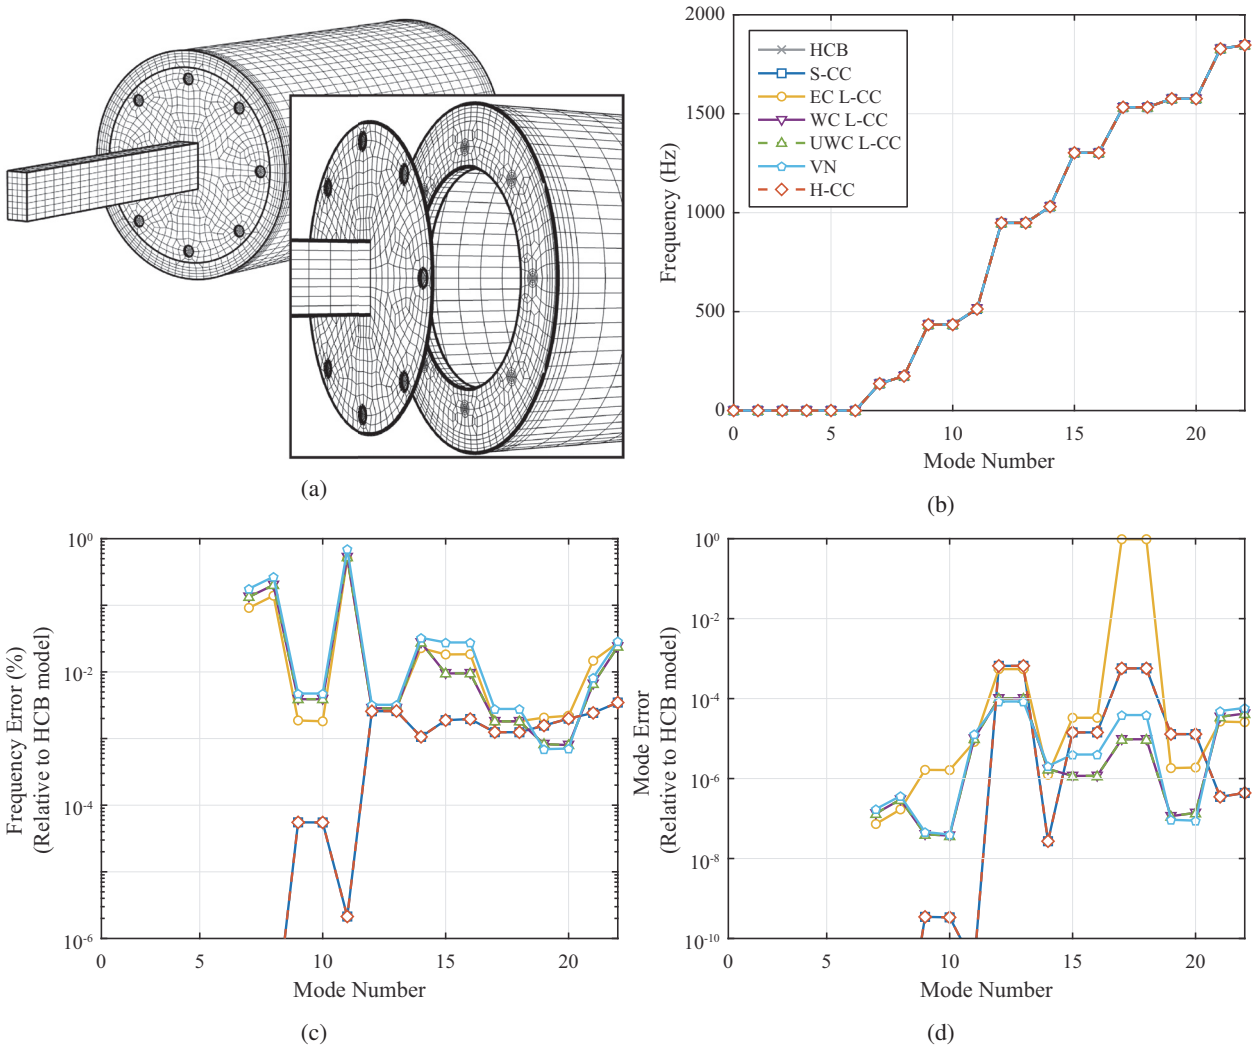

Figure 10: FE mesh (a), natural-frequency comparison (b), frequency-error comparison (c), and mode-error comparison (d) for the can-beam model.

Figures 10c and 10d show the error in natural frequencies and mode shapes respectively. Even though the interface between the two substructures is clearly divided into 8 patches, from an analysis standpoint it is treated as a single interface. This means that just like for the cantilevered plate, the S-CC results are identical to the H-CC results, and the WC L-CC results are identical to the UWC L-CC results. The frequency and mode error plots behave quite a bit differently than in the previous examples. This is expected because the can-beam model has small, concentrated interfaces that are relatively rigid compared to the broad, flexible interfaces in the previous examples. As a result, the VN method performs much better for this type of interface because a single VN can quite easily capture the dynamics of one of the contact patches. As long as at least 8 VNs are kept to represent the bolted interface, the reduction performs quite well. The EC L-CC and WC L-CC / UWC L-CC methods perform slightly better still, and once again the best performance comes from the S-CC / H-CC methods. Modes 12 and 13, and modes 17 and 18 show degeneracy in the mode shape error due to the symmetry of the can substructure. This explains the abnormally high mode error for these two mode sets across all methods. The presence of this degeneracy also suggests that these modes (and potentially others as well) have displacements that are localized within the can substructure and thus have very little interface participation. Low interface participation may explain why some modes (e.g., 9, 10, 12, 13, 17, and 18) have much lower frequency error than others.

Table 10: The number (and cutoff frequency where applicable) of interface modes used by each method in reducing the can-beam model

Method	Number of CC Modes Calculated Per Substructure (Maximum Frequency, Hz)		Number of Interface DOF (Maximum Frequency, Hz)
	Sub. 1	Sub. 2	
HCB	-	-	2232
S-CC / H-CC	-	-	48 (38880)
EC L-CC	27 (12240)	27 (12930)	48
WC L-CC / UWC L-CC	72 (100040)	72 (87780)	48
VN	-	-	8×6 = 48

Table 11: Computation times needed to create each reduced-order model for the can-beam model, and the computation time needed to solve for each reduce-order model's eigenvalues. (Note that the reduced-order modeling times for each interface reduction only include the time needed to compute the interface reduction and do not include the HCB modeling time.)

Method	Model Reduction Time (ms)	Eigenvalue Solution Time (ms)
Full FE	-	56,462
HCB	201,865	3,729
S-CC	5,805	18
H-CC	7,630	9
EC L-CC	11,444	25
WC L-CC	32,334	37
UWC L-CC	31,906	20
VN	937	10

The computation times for the can-beam model are summarized in Table 11. The model reduction time reported for each interface reduction is the time needed just to perform the interface reduction, and does not include the time needed to obtain the HCB model. Note that this model shows computation times for an industrial scale model, but the relative computation times could still vary significantly for systems with more than two substructures, and more complex interface configurations.

A performance study is carried out for the can-beam model by incrementing the cut off frequency and obtaining the resulting maximum frequency error over the first 22 modes versus the number of interface DOF. Performance curves are obtained in this way for each interface reduction method and plotted in Fig. 11. The local methods all have quite similar performance. They achieve good frequency error (<1%) very quickly, but then plateau and do not not fall below 0.1% error without exceeding 150 interface DOF. As discussed above, the VN method performs much better for the can-beam system and does nearly as well as the other local methods. Table 12 summarizes the number of modes necessary for each method to reach a frequency error of less than 1%.

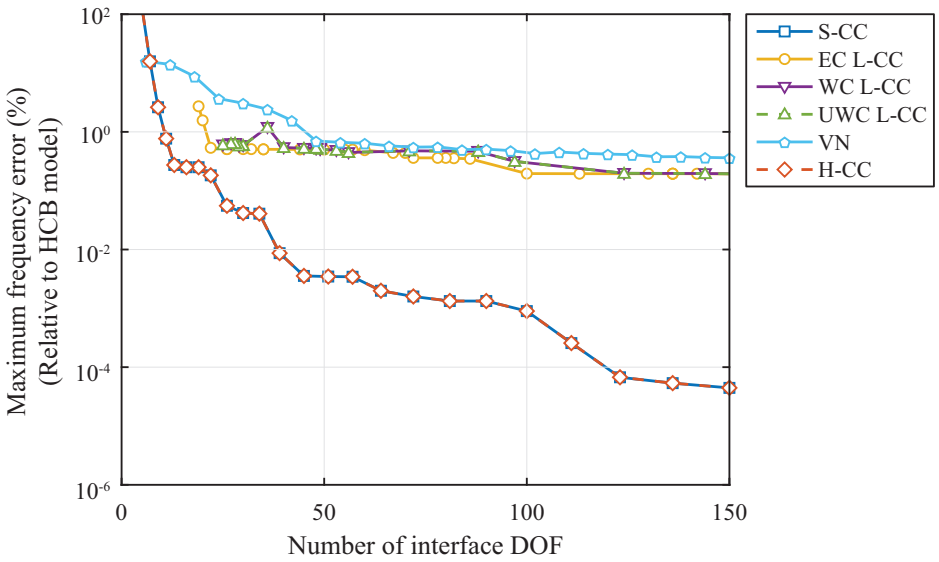

Figure 11: Study of maximum frequency error with increasing number of interface DOF for various interface reduction methods using the can-beam model

Table 12: Minimum interface size to achieve 1% frequency error for the can-beam model

	S-CC / H-CC	EC L-CC	WC L-CC / UWC L-CC	VN
Number of CC Modes	11	22	25	48
Maximum Frequency, Hz	1810	3780	28130	-

4. Conclusion

This work provides a review and improvements on six different interface reduction techniques that can be applied to Hurty/Craig-Bampton models and provides comparative examples to demonstrate the advantages and disadvantages of each modeling approach:

1. System-level characteristic constraint
2. Exact-compatibility local-level characteristic constraint
3. Weak-compatibility local-level characteristic constraint
4. Uncoupled weak-compatibility local-level characteristic constraint
5. Virtual node
6. Hybrid-level characteristic constraint

Each method can be classified as either system-level (item 1), local-level (items 2-5) or hybrid-level (item 6). The system-level technique applies the interface reduction using information from the fully assembled Hurty/Craig-Bampton models whereas any of the four local-level techniques can be applied to individual substructures without explicitly knowing the connecting substructures. The hybrid-level reduction falls in between these two since the reduction can be applied locally to a single interface but still requires information about the immediately adjacent substructures.

Modal analysis solutions obtained by the system-level method are several orders of magnitudes more accurate when compared with the local-level methods because the coupling between all substructures has been properly considered. However, the computational cost of the system-level eigenvalue analysis may be more expensive or even unacceptable for large-scale FE models with millions of interface DOF. Under these conditions, the local methods, where the modal analyses are applied at the substructure level before the system is assembled, are preferred. Local

methods also retain a greater degree of versatility because they can be assembled in different configurations after performing the reduction. Three of the four local-level reductions utilize interface vibration modes with either exact or weak interface compatibility, and the last uses virtual nodes to reduce the interface. The three based on interface vibration modes have comparable performance, and the virtual node method consistently has the worst performance when applied to models with flexible interfaces. The hybrid-level technique achieves accuracy comparable to that of the system-level technique because it utilizes assembled interfaces to produce the reduced-order basis, but also achieves some of the benefits of local-level reduction. By computing the interface modes for sub-assemblies rather than the whole system assembly, the computational cost of obtaining hybrid-level interface modes may be reduced compared to the system-level interface modes, especially when there are many substructures with many interfaces. Furthermore, when substructures are added, removed, or modified, only those interfaces connecting to the updated substructures need to be recalculated.

The interface reductions presented in this paper are all based on methods from the literature, but several have been improved or generalized in some way. Both of the local-level weak-compatibility methods are generalizations of the method presented in [17] to systems with multiple interfaces. The hybrid-level technique is similar to the one presented in [2], but the theory has been revised to allow for the rigid-body motion to be captured in the formulation. This modification allows the method to be used for systems with free boundary conditions, resulting in a more accurate set of predictions.

There are several open areas for future research relating to the interface reduction methods presented in this paper. Some or all of the methods may be extended to nonlinear substructure models. Nonlinear calculations tend to be even more computationally demanding than linear models, so interface reduction could have a dramatic impact in reducing computational cost. Another possibility is to extend the interface reduction techniques to different CMS representations such as those based on free-interface and attachment modes, or those that extend to higher frequencies by using quasi-static constraint modes. Finally, it is possible to significantly enhance the accuracy of HCB models by approximating the participation of the residual modes [15]. The same approach can be used to enhance the accuracy of system-level interface reduction, and may also be extended to the local or hybrid-level techniques.

Acknowledgements

The authors would like to acknowledge Pascal Reuss and Loic Salles for their discussions and input during the early stages of this work. Part of this research was conducted at the 2016 Nonlinear Mechanics and Dynamics (NOMAD) Research Institute sponsored by Sandia National Laboratories and hosted at the University of New Mexico. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

[1] Allemang, B J. ; Drown, D L.: A Correlation Coefficient for Modal Vector Analysis. In: *Proceedings, International Modal Analysis Conference*, 1982, S. 110–116

[2] Aoyama, Y. ; Yagawa, G.: Component mode synthesis for large-scale structural eigenanalysis. In: *Computers & Structures* 79 (2001), Nr. 6, S. 605–615

[3] Balmès, E.: Use of Generalized Interface Degrees of Freedom in Component Mode Synthesis. In: *Proceedings of the 14th International Modal Analysis Conference* 1 (1996)

[4] Bampton, M. C. C. ; Craig, Jr R. R.: Coupling of substructures for dynamic analyses. In: *AIAA Journal* 6 (1968), Nr. 7, S. 1313–1319

[5] Castanier, M. P. ; Tan, Y.-C. ; Pierre, C.: Characteristic constraint modes for component mode synthesis. In: *AIAA journal* 39 (2001), Nr. 6, S. 1182–1187. – ISSN 0001–1452

[6] Craig, Roy R.: Coupling of substructures for dynamic analyses: an overview. In: *Proceedings of AIAA/ASME/ASCE/AHS/ASC structures, structural dynamics, and materials conference and exhibit*, 2000, S. 1573–1584

[7] Craig, Roy R. ; Kurdila, Andrew J.: *Fundamentals of structural dynamics*. John Wiley & Sons, 2006

[8] Craig Jr, Roy R. ; Chang, Ching-jone: On the use of attachment modes in substructure coupling for dynamic analysis. In: *transformation* 10 (1977), Nr. 4, S. 3

[9] Craig Jr, Roy R. ; Chang, Ching-Jone: Substructure coupling for dynamic analysis and testing / National Aeronautics and Space Administration. 1977. – Report

[10] Herrmann, Jan ; Maess, Matthias ; Gaul, Lothar: Substructuring including interface reduction for the efficient vibro-acoustic simulation of fluid-filled piping systems. In: *Mechanical Systems and Signal Processing* 24 (2010), Nr. 1, S. 153–163. – ISSN 0888–3270

[11] Holzwarth, Philip ; Eberhard, Peter: Interface Reduction for CMS Methods and Alternative Model Order Reduction. In: *IFAC-PapersOnLine* 48 (2015), Nr. 1, S. 254–259

[12] Hong, Sung-Kwon ; Epureanu, Bogdan I. ; Castanier, Matthew P.: Next-generation parametric reduced-order models. In: *Mechanical Systems and Signal Processing* 37 (2013), Nr. 1–2, S. 403–421

[13] Hurty, Walter C.: Vibrations of structural systems by component mode synthesis. In: *Journal of the Engineering Mechanics Division* 86 (1960), Nr. 4, S. 51–70
[14] Hurty, Walter C.: Dynamic analysis of structural systems using component modes. In: *AIAA journal* 3 (1965), Nr. 4, S. 678–685
[15] Kim, Jin-Gyun ; Lee, Phill-Seung: An enhanced Craig–Bampton method. In: *International Journal for Numerical Methods in Engineering* 103 (2015), Nr. 2, S. 79–93
[16] Klerk, D. D. ; Rixen, D. J. ; Voormeeren, S. N.: General Framework for Dynamic Substructuring: History, Review and Classification of Techniques. In: *AIAA Journal* 46 (2008), Nr. 5, S. 1169–1181. – ISSN 0001–1452
[17] Kuether, Robert J. ; Allen, Matthew S. ; Hollkamp, Joseph J.: Modal Substructuring of Geometrically Nonlinear Finite Element Models with Interface Reduction. In: *AIAA Journal* 55 (2017), Nr. 5, S. 1695–1706. – ISSN 0001–1452
[18] Leung, Andrew Yee-Tak: An accurate method of dynamic condensation in structural analysis. In: *International Journal for Numerical Methods in Engineering* 12 (1978), Nr. 11, S. 1705–1715
[19] Lindberg, Eskil ; Hörlin, Nils-Erik ; Göransson, Peter: Component Mode Synthesis Using Undeformed Interface Coupling Modes to Connect Soft and Stiff Substructures. In: *Shock and Vibration* 20 (2013), S. 157–170
[20] MacNeal, Richard H.: Special Issue on Structural Dynamics A hybrid method of component mode synthesis. In: *Computers & Structures* 1 (1971), Nr. 4, S. 581–601
[21] Rixen, Daniel J.: Force modes for reducing the interface between substructures. In: *Proceedings of IMAC-XX: A Conference on Structural Dynamics, Society for Experimental Mechanics, Bethel, CT* Bd. 2, S. 8
[22] Rixen, Daniel J.: A dual Craig–Bampton method for dynamic substructuring. In: *Journal of Computational and Applied Mathematics* 168 (2004), Nr. 1–2, S. 383–391. – ISSN 0377–0427
[23] Rubin, S.: Improved Component-Mode Representation for Structural Dynamic Analysis. In: *AIAA Journal* 13 (1975), Nr. 8, S. 995–1006
[24] Shyu, Wen-Hwa ; Gu, Jianmin ; Hulbert, Gregory M. ; Ma, Zheng-Dong: On the use of multiple quasi-static mode compensation sets for component mode synthesis of complex structures. In: *Finite Elements in Analysis and Design* 35 (2000), Nr. 2, S. 119–140. – ISSN 0168–874X
[25] Shyu, Wen-Hwa ; Ma, Zheng-Dong ; Hulbert, Gregory M.: A new component mode synthesis method: Quasi-static mode compensation. In: *Finite Elements in Analysis and Design* 24 (1997), Nr. 4, S. 271–281. – ISSN 0168–874X
[26] Tran, D. M.: Component mode synthesis methods using interface modes. Application to structures with cyclic symmetry. In: *Computers & Structures* 79 (2001), Nr. 2, S. 209–222. – ISSN 0045–7949
[27] Tran, Duc-Minh: Component mode synthesis methods using partial interface modes: Application to tuned and mistuned structures with cyclic symmetry. In: *Computers & Structures* 87 (2009), Nr. 17–18, S. 1141–1153

Appendices

Appendix A User Selected Parameters

Several of the interface reduction methods described in the main body of this paper require the user to choose some cut off parameters. This section provides a few studies designed to show the effect that these parameters have on the performance (i.e. accuracy versus model size) of each method. The studies were performed on the W-bracket model and thus the parameters selected as a result may not be optimal in general, however the results display the general effect that the parameters have on the performance.

A.1 EC L-CC

When implementing the EC L-CC interface reduction, the user must decide on a range-space truncation tolerance, σ_R, to be used in conjunction with the SVD orthogonalization in Eq. (18). Selecting a truncation tolerance $\sigma_R = 0$ will result in a pure orthogonalization with no truncation. This is practically equivalent to a Graham-Schmidt or QR orthogonalization. As σ_R is increased, the level of truncation grows, thus reducing the model size and increasing error. To find the value of σ_R that works best for the W-bracket model, a performance study was carried out. For each value of σ_R, the number of interface DOF in the EC L-CC model was slowly increased by increasing the frequency cutoff used to obtain the interface mode sets. For each of these models, the maximum frequency error over the first 26 elastic modes was evaluated with respect to the HCB model. The performance curves are plotted in Fig. A.1. Overall, the best performance was achieved using $\sigma_R = 1 \times 10^{-4}$. This agrees with the choice of σ_R used in [12].

A.2 WC L-CC

To implement the WC L-CC interface reduction, the user must decide on an approximate null-space tolerance, σ_N, used to weakly enforce the constraint equations. In the WC L-CC method, error occurs due to (1) interface mode truncation, and (2) inexact compatibility. Raising σ_N will increase the number of interface DOF and thus decrease

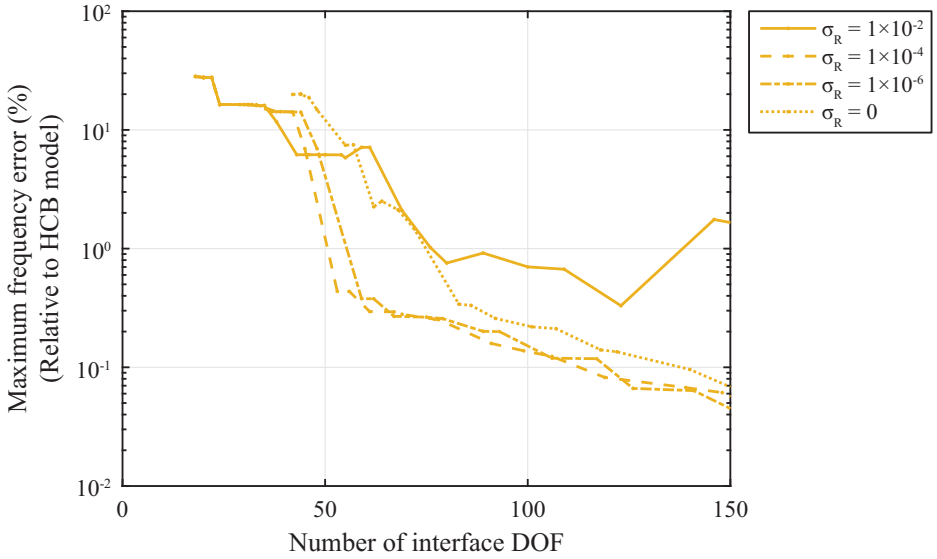

Figure A.1: EC L-CC performance curves showing maximum frequency error versus number of interface DOF for the W-bracket model. Different curves are obtained by varying the truncation tolerance, σ_R.

the first kind of error, but will relax the compatibility tolerance and thus increase the second kind of error. To find the value of σ_N that works best for the W-bracket model, a performance study was carried out. For each value of σ_N, the number of interface DOF in the WC L-CC model was incremented and the maximum frequency error over the first 26 elastic modes was evaluated with respect to the HCB model. The performance curves are plotted in Fig. A.2. The best performance was achieved using $\sigma_N = 1 \times 10^{-1}$. It is interesting that optimal results were obtained with such a low value of σ_N; apparently this model is more sensitive to mode truncation than inexact compatibility.

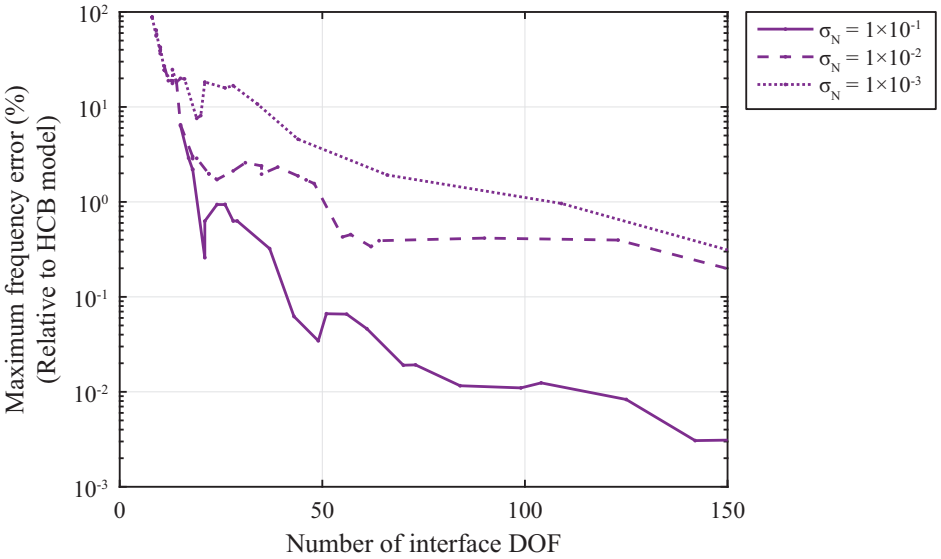

Figure A.2: WC L-CC performance curves showing maximum frequency error versus number of interface DOF for the W-bracket model. Different curves were obtained by varying the null-space tolerance, σ_N.

A.3 UWC L-CC

Similar to the WC L-CC method, the WC L-CC method requires a null-space tolerance, σ_N. To find the best value of σ_N, a performance study was carried out exactly the same way as in the previous section. The resulting performance curves are shown in Fig. A.3. Once again, the best performance was observed with $\sigma_N = 1 \times 10^{-1}$.

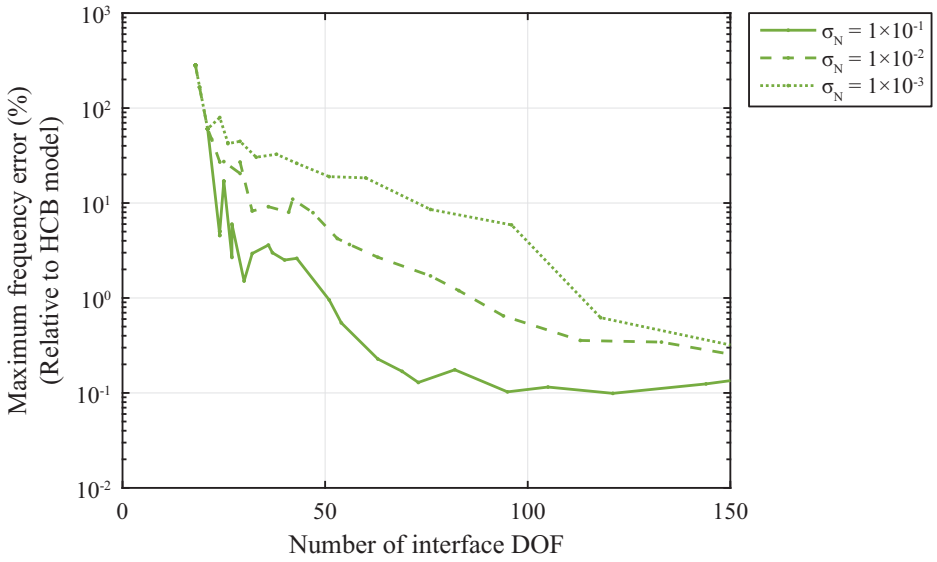

Figure A.3: UWC L-CC performance curves showing maximum frequency error versus number of interface DOF for the W-bracket model. Different curves were obtained by varying the null-space tolerance, as well as the mode selection method.