JOURNAL OF GUIDANCE, CONTROL, AND DYNAMICS
Vol. 20, No. 6, November–December 1997

Orthogonal Square Root Eigenfactor Parameterization of Mass Matrices

John L. Junkins* and Hanspeter Schaub†
Texas A&M University, College Station, Texas 77843-3141

An improved method is presented to parameterize a smoothly time varying, symmetric, positive definite system mass matrix $M(t)$ in terms of the instantaneous eigenfactors, namely, the eigenvalues and eigenvectors of $M(t)$. Differential equations are desired whose solutions generate the instantaneous spectral decomposition of $M(t)$. The derivation makes use of the fact that the eigenvector matrix is orthogonal and, thus, evolves analogously to a higher-dimensional rotation matrix. Careful attention is given to cases where some eigenvalues and/or their derivatives are equal or near equal. A robust method is presented to approximate the corresponding eigenvector derivatives in these cases, which ensures that the resulting eigenvectors still diagonalize the instantaneous $M(t)$ matrix. This method is also capable of handling the rare case of discontinuous eigenvectors, which may only occur if both the corresponding eigenvalues and their derivatives are equal.

I. Introduction

MOST multibody dynamical systems such as multilink robots have configuration-dependent mass matrices. This dependency makes the mass matrix vary with time. Solving such dynamical systems involves performing an inverse of the system mass matrix at each integration step. Finding this inverse is computationally difficult and expensive for large systems. Furthermore, standard inverse and linear equations solution techniques are not easily parallelizable and, therefore, cannot take full advantage of modern parallel computing systems.

A method is introduced that parameterizes the symmetric, positive definite mass matrix in terms of its eigenfactors, i.e., eigenvectors and eigenvalues. Instead of forward integrating the original mass matrix differential equation directly, only the eigenfactors themselves are forward integrated. The resulting formulation is one that could be easily implemented on a massively parallel computer system. A paper by Oshman and Bar-Itzhack[1] introduces an important orthogonal square root eigenfactor parameterization to solve the differential matrix Riccati equation. However, some details of their treatment of equal or near-equal eigenvalues were found to be incorrect, and the case of discontinuous eigenvectors was not accounted for. This paper provides an approximate treatment of the near-equal eigenvalue case and an associated error analysis where the symmetric, positive definite matrix parameterized is a state-dependent mass matrix $M(\boldsymbol{x}, t)$.

Eigenfactor derivatives have been discussed in the literature for quite some time, but they are mostly used to establish modal sensitivities and not to derive eigenfactor differential equations.[2,3] A majority of the engineering literature on eigenfactor derivatives is for the general structural eigenvalue problem $(K - \lambda_i M)\boldsymbol{v}_i = 0$. This paper deals with the problem where we need to find the eigenfactors and their derivatives for given continuous matrices $M(\boldsymbol{x}, t)$ and $\dot{M}(\boldsymbol{x}, \dot{\boldsymbol{x}}, t)$. Inasmuch as M is symmetric, its eigenvector matrix is guaranteed to be orthogonal and, thus, behave analogously to a higher-dimensional rotation direction cosine matrix.[4,5] This analogy is the starting point for a simple derivation of the eigenfactor derivatives. Except for repeated eigenvalues, for smooth $M(\boldsymbol{x}, t)$, we anticipate continuous differentiable eigenfactors.

When repeated eigenvalues are present, there are often numerical problems associated with calculating the eigenvector derivatives. The literature usually deals with the case where M and $\dot{M}$ are given and one needs to solve for the eigenfactor derivatives.[6–9] A simple method is presented that usually allows the eigenvectors to be smoothly integrated from a differential equation, even including the case of repeated eigenvalues, thus avoiding having to solve an algebraic eigenvalue problem again.

Repeated eigenvalues for mass matrices are common in mechanics, e.g., principal axes for mass matrices and stress tensors. We note that, if the objective is to find a diagonalizing transformation, then near-repeated eigenvalues only indicate a loss of uniqueness of the eigenvectors. For many purposes uniqueness is not required, and any set of orthogonal eigenvectors that span a unique subspace is admissible. This is the approach realized in the present paper when eigenvectors lose their uniqueness for near-repeated eigenvalues; we seek to generate admissible orthogonal eigenvectors that span the correct subspace to within acceptable precision. This line of thinking is easily implemented for solving the algebraic eigenvalue problem for a given constant matrix; however, generating the time varying instantaneous eigensolution by solving differential equations governing the eigenfactors is a more challenging task that is addressed.

II. Problem Statement

Any real, symmetric, positive definite matrix M of dimension $n \times n$ can be decomposed into n positive real eigenvalues λ_i and n orthogonal eigenvectors $\boldsymbol{c}_i$:

$$M = C^T \Lambda C \tag{1}$$

where $\Lambda = \text{diag}(\lambda_i)$ and C is defined as

$$C = [\boldsymbol{c}_1, \ldots, \boldsymbol{c}_n]^T \tag{2}$$

The evolution of the C eigenvector matrix is completely analogous to the time variation of an orthogonal direction cosine attitude matrix. Because M is symmetric and positive definite, the eigenvector matrix C is guaranteed to be orthogonal:

$$C^T C = C C^T = I \tag{3}$$

and the eigenvalues will always be positive. Thus, we can always define $s_i = \sqrt{\lambda_i}$ and the corresponding matrix $S = \text{diag}(s_i)$. This allows M to be written as

$$M = C^T S^T S C = W^T W \tag{4}$$

where $W = SC$ is a matrix square root of M. Therefore, keeping track of C and S is equivalent to a square root algorithm with all of

Received Nov. 10, 1996; presented as Paper 97-0248 at the AIAA 35th Aerospace Sciences Meeting, Reno, NV, Jan. 6–9, 1997; revision received May 29, 1997; accepted for publication June 10, 1997.

*George J. Eppright Chair Professor, Department of Aerospace Engineering. Fellow AIAA.

†Graduate Research Assistant, Department of Aerospace Engineering. Student Member AIAA.

the associated numerical robustness advantages.[1] Note that, because $s_i, \lambda_i > 0$, the matrix inverse of M is trivial:

$$M^{-1} = C^T S^{-2} C \tag{5}$$

Let $\boldsymbol{x}$ be a continuous system state vector satisfying the second-order dynamical differential equation

$$M(\boldsymbol{x})\ddot{\boldsymbol{x}} = F(\boldsymbol{x}, \dot{\boldsymbol{x}}, \boldsymbol{u}) \tag{6}$$

where $\boldsymbol{u}$ is a control vector. Any discontinuity in $\boldsymbol{u}$ can only directly affect $\ddot{\boldsymbol{x}}$ and not $\dot{\boldsymbol{x}}$ or $\boldsymbol{x}$. Therefore, the matrices $M(\boldsymbol{x})$ and $\dot{M}(\boldsymbol{x}, \dot{\boldsymbol{x}}, t)$ are continuous. Note that complete algebraic expressions for $M(\boldsymbol{x})$ and $\dot{M}(\boldsymbol{x}, \dot{\boldsymbol{x}}, t)$ are usually available for dynamical systems.

III. Eigenfactor Derivatives

The fact that the eigenvector matrix C that parameterizes a symmetric matrix M is always orthogonal will be central in the following derivation of the eigenfactor derivatives. It is known that a smoothly varying $n \times n$ orthogonal matrix $C(t)$ satisfies the same Poisson differential equation as does the attitude direction cosine matrix[5,10,11]

$$\dot{C} = -[\Omega]C \tag{7}$$

where $[\Omega]$ is a skew-symmetric matrix,

$$[\Omega] = \begin{bmatrix} 0 & \Omega_{12} & \Omega_{13} \\ -\Omega_{12} & 0 & \Omega_{23} \\ -\Omega_{13} & -\Omega_{23} & 0 \end{bmatrix} \tag{8}$$

For the case of attitude direction cosine matrices, these Ω_{ij} terms represent body angular velocities. This concept of angular velocities lifts to the case of higher-dimensional orthogonal matrices. Instead of being body angular velocities, here the Ω_{ij} terms can be viewed as eigenvector angular velocities. Each Ω_{ij} term gives a measure of the rate that the two eigenvectors $\boldsymbol{c}_i$ and $\boldsymbol{c}_j$ are rotating in the plane spanned by these two vectors.

The problem of finding eigenvector derivatives has now been reduced to one of finding the Ω_{ij} terms. Taking the first derivative of Eq. (1), we obtain

$$\dot{M} = C^T \Lambda \dot{C} + \dot{C}^T \Lambda C + C^T \dot{\Lambda} C \tag{9}$$

After making use of Eq. (7) and defining the symmetric ξ matrix to be

$$\xi = [\Omega]\Lambda - \Lambda[\Omega] \tag{10}$$

Eq. (9) is reduced to

$$\dot{M} = C^T(\xi + \dot{\Lambda})C \tag{11}$$

The ξ matrix can be shown to be

$$\xi_{ij} = \Omega_{ij}(\lambda_j - \lambda_i) \tag{12}$$

Note that, because the diagonal terms of ξ are zero and $\dot{\Lambda}$ is a diagonal matrix, these two matrices projected through C split up the $\dot{M}$ matrix into diagonal and off-diagonal terms. It turns out that the diagonal terms lead to the eigenvalue derivatives and the off-diagonal terms leads to the angular velocity terms required to determine the eigenvector derivatives from Eq. (7). Let the matrix μ be

$$\mu = C\dot{M}C^T \tag{13}$$

Using Eq. (12), this is rewritten as

$$\mu = \xi + \dot{\Lambda} \tag{14}$$

which can be component-wise expressed as

$$\mu_{ij} = \begin{cases} \Omega_{ij}(\lambda_j - \lambda_i) & \text{for} \quad i \neq j \\ \dot{\lambda}_i & \text{for} \quad i = j \end{cases} \tag{15}$$

Note that Eq. (15) must always hold, even in the presence of repeated eigenvalues. Also, as long as Ω_{ij} is bounded, then μ_{ij} must go to zero as $\lambda_i \to \lambda_j$. From here it is trivial to express Ω_{ij} as[1,2,4]

$$\Omega_{ij} = \frac{\mu_{ij}}{\lambda_j - \lambda_i} \qquad \text{for} \qquad \lambda_i \neq \lambda_j \tag{16}$$

or in terms of the square roots of the eigenvalues as

$$\Omega_{ij} = \frac{\mu_{ij}}{s_j^2 - s_i^2} \qquad \text{for} \qquad s_i \neq s_j \tag{17}$$

If M is a system mass matrix, then it is positive definite and all of its eigenvalues are positive. Therefore, the square roots of the eigenvalues will always be defined and real. At first glance it might appear that Ω_{ij} will go to infinity when $\lambda_i \to \lambda_j$. However, this is generally not the case, as will be shown in the next section.

Had an eigenvector matrix $V = C^T$ been used instead to parameterize M in Eq. (1), then V would have also been orthogonal and abided by the same differential equation as in Eq. (7) with a different $[\Omega]$ matrix,

$$\dot{V} = -[\bar{\Omega}]V \tag{18}$$

However, finding this $[\bar{\Omega}]$ matrix is more difficult and can only be achieved through a coordinate transformation of the $[\Omega]$ matrix

$$[\bar{\Omega}] = -V[\Omega]V^T \tag{19}$$

In rotational dynamics, this would be analogous to writing the body angular velocity matrix in the inertial frame. This is why the current development uses C instead of V.

Expressions for $\dot{\lambda}_i$ are found directly from Eq. (15):

$$\dot{\lambda}_i = \mu_{ii} = \Gamma_{ii} \tag{20}$$

where $\Gamma = \text{diag}(\mu_{ii})$. Because $s_i = \sqrt{\lambda_i}$, the derivative of s_i is

$$\dot{s}_i = \tfrac{1}{2}\dot{\lambda}_i \lambda_i^{-\frac{1}{2}} = \tfrac{1}{2}\dot{\lambda}_i s_i^{-1} \tag{21}$$

or in matrix form[1,4]

$$\dot{S} = \tfrac{1}{2}\Gamma S^{-1} \tag{22}$$

Note that calculating s_i^{-1} is always possible because for symmetric, positive definite matrices all eigenvalues are always positive.

It turns out that by far the largest computational effort in evaluating $\dot{C}$ and $\dot{S}$ is finding the $\mu_{ij} = c_i^T \dot{M} c_j$ terms. However, this inner product operation lends itself perfectly to be performed on a massively parallel computer system where all of the μ_{ij} terms could be calculated independently in parallel. By contrast, direct matrix inverse methods do not lend themselves easily to be calculated on parallel systems.

IV. Repeated Eigenvalues

A. Distinct Eigenvalue Derivatives

Clearly, mathematical problems arise computing Ω_{ij} using Eq. (16) or (17) when $\lambda_i \to \lambda_j$. Whenever $\lambda_i = \lambda_j$, the associated eigenvectors $\boldsymbol{c}_i$ and $\boldsymbol{c}_j$ that parameterize M are not unique. Any two orthogonal unit vectors in the plane spanned by $\boldsymbol{c}_i$ and $\boldsymbol{c}_j$ would generate the proper M matrix. However, the eigenfactors not only have to generate the proper M matrix, but they and their derivatives also need to generate the proper continuous $\dot{M}$ matrix. This leads to the following proposition about the continuity of $\boldsymbol{c}_i$ and $\boldsymbol{c}_j$ of any continuous matrix M through the point where the eigenvalues λ_i and λ_j are equal.

Proposition 1: Let M be a continuous matrix with a continuous derivative $\dot{M}$, and let λ_i and λ_j be repeated eigenvalues of M with distinct derivatives $\dot{\lambda}_i \neq \dot{\lambda}_j$. Then the associated eigenvectors will be continuous and unique through the condition where $\lambda_i = \lambda_j$.

Proof: To prove this proposition, let us write M as[12]

$$M = \sum_{i=1}^{n} \lambda_i E_i \tag{23}$$

where E_i is defined as the outer product

$$E_i = \boldsymbol{c}_i \boldsymbol{c}_i^T \tag{24}$$

The matrix $\dot{M}$ can now be written as

$$\dot{M} = \sum_{i=1}^{n} (\dot{\lambda}_i E_i + \lambda_i \dot{E}_i) \tag{25}$$

Without loss of generality, we can assume that only the first two eigenvalues λ_1 and λ_2 are repeated. Let us define E_{ij} as

$$E_{ij} = E_i + E_j \tag{26}$$

then we can rewrite $\dot{M}$ as

$$\dot{M} = \dot{\lambda}_1 E_1 + \dot{\lambda}_2 E_2 + \lambda_1 \dot{E}_{12} + \sum_{i=3}^{n} (\dot{\lambda}_i E_i + \lambda_i \dot{E}_i) \tag{27}$$

Because $\dot{M}$ must be continuous, $\boldsymbol{c}_1$ and $\boldsymbol{c}_2$ cannot be discontinuous because the eigenvalue derivatives $\dot{\lambda}_1$ and $\dot{\lambda}_2$ are distinct. □

Proposition 1 leads to the following two corollaries regarding the boundedness of Ω_{ij} for the case where the matrix M is symmetric.

Corollary 1: Assume M and $\dot{M}$ are continuous and symmetric. Whenever $\lambda_i \to \lambda_j$ and $\dot{\lambda}_i \not\to \dot{\lambda}_j$, then $\Omega_{ij} \not\to \pm\infty$.

Proof: If Ω_{ij} would go to $\pm\infty$ as $\lambda_i \to \lambda_j$ in Eq. (16) or (17), then it would be impossible for $\boldsymbol{c}_i$ and $\boldsymbol{c}_j$ to be continuous through the point where $\lambda_i = \lambda_j$, which would contradict Proposition 1. □

Corollary 1 shows that an assumption made in Ref. 1 is generally not true. There it is assumed that, whenever $\lambda_i \to \lambda_j$, $\Omega_{ij} \to \pm\infty$. The magnitude of Ω_{ij} is used to check whether two eigenvalues are very close. However, for the case of repeated eigenvalues having distinct eigenvalue derivatives, Corollary 1 shows that this is not the case.

Corollary 2: Assume M and $\dot{M}$ are continuous and symmetric. Whenever $\lambda_i \to \lambda_j$ and $\dot{\lambda}_i \neq \dot{\lambda}_j$, then $\mu_{ij} \to 0$.

Proof: Because Eq. (15) must hold for all time and Ω_{ij} must be bounded whenever $\lambda_i \to \lambda_j$ and $\dot{\lambda}_i \not\to \dot{\lambda}_j$, then μ_{ij} must go to zero. □

From Corollaries 1 and 2 it is clear that any difficulties in calculating Ω_{ij} for repeated eigenvalues with distinct eigenvalue derivatives are purely numerical problems. Numerical simulations indicate that it is usually easy to calculate Ω_{ij} in the neighborhood of two repeated eigenvalues as long as the values in the matrix $\dot{M}$ are not excessively large. In fact, Mills-Curran[7] shows that, at the instant where $\lambda_i = \lambda_j$ and $\dot{\lambda}_i \neq \dot{\lambda}_j$, a term analogous to Ω_{ij} can be calculated using the distinct eigenvalue derivatives:

$$\Omega_{ij} = \frac{\boldsymbol{c}_i^T \ddot{M} \boldsymbol{c}_j}{2(\dot{\lambda}_j - \dot{\lambda}_i)} \tag{28}$$

However, this expression is not valid in the neighborhood of repeated eigenvalues. In a numerical simulation one usually does not have exact repeated eigenvalues. To avoid numerical problems of calculating the angular velocity terms in the neighborhood of repeated eigenvalues, an approximation is introduced.[1] Assume Ω_{ij} is known at time step t_0. Then Ω_{ij} at time step $t_1 > t_0$ is approximated as

$$\Omega_{ij}(t_1) = \begin{cases} \dfrac{\mu_{ij}}{\lambda_j - \lambda_i} & \text{for} \quad |\lambda_i - \lambda_j| \geq \epsilon \\ \Omega_{ij}(t_0) + \dot{\Omega}_{ij}(t_0)(t_1 - t_0) & \text{for} \quad |\lambda_i - \lambda_j| < \epsilon \end{cases} \tag{29}$$

In other words, the Ω_{ij} term is approximated linearly for as long as $|\lambda_i - \lambda_j| < \epsilon$. This approximation is valid for any continuous, symmetric matrix M with a continuous derivative. The derivative $\dot{\Omega}_{ij}(t_0)$ can be found numerically through a backward difference method. For the present discussion we ignore the error in approximating $\dot{\Omega}_{ij}$.

Proposition 2: Let the matrices M and $\dot{M}$ be continuous and symmetric. Assume that $|\lambda_i - \lambda_j| < \epsilon$ and that $|\dot{\lambda}_i - \dot{\lambda}_j| \gg \epsilon$ over the time span $[t_1, t_2]$; then the Ω_{ij} approximation in Eq. (29) will introduce an error in $\dot{M}$ of the order of ϵ^3.

Proof: Without loss of generality, let us assume that only the λ_1 and λ_2 are within ϵ to being repeated eigenvalues:

$$\lambda_2 = \lambda_1 + \epsilon \tag{30}$$

Let $\dot{M}$ be the exact derivative obtained from the true $[\Omega_{ij}]$ matrix and $\underline{\dot{M}}$ be the approximated derivative due to the approximated $\underline{\Omega}_{ij}$ terms using Eq. (29). The error in $\dot{M}$ is given as

$$\Delta\dot{M} = \underline{\dot{M}}(\Lambda, C, [\underline{\Omega}]) - \dot{M}(\Lambda, C, [\Omega]) \tag{31}$$

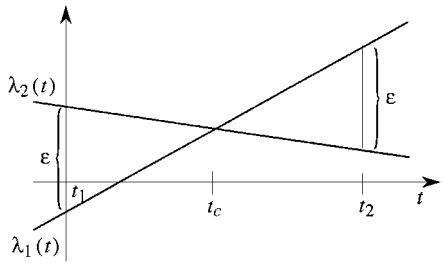

Fig. 1 Repeated eigenvalue illustration.

After using Eqs. (7), (24), and (25), $\Delta\dot{M}$ can be expressed as

$$\Delta\dot{M} = \epsilon \Delta\Omega_{12} \left(\boldsymbol{c}_1 \boldsymbol{c}_2^T + \boldsymbol{c}_2 \boldsymbol{c}_1^T \right) \tag{32}$$

The time span $\Delta t = t_2 - t_1$, where $|\lambda_1 - \lambda_2| < \epsilon$, can be estimated as follows. Let t_1 be the time where the eigenvalue difference is less than ϵ and t_2 the time where the difference grows larger again than ϵ. If it is assumed that $|\dot{\lambda}_1 - \dot{\lambda}_2| \gg \epsilon$, then λ_1 and λ_2 can be assumed to be near linear between t_1 and t_2 as shown in Fig. 1. The time Δt is approximately related to the difference in eigenvalues ϵ through

$$\Delta t \simeq \frac{2\epsilon}{\dot{\lambda}_1 - \dot{\lambda}_2} \tag{33}$$

During the interval $[t_1, t_2]$, $\underline{\Omega}_{12}$ is linearly extrapolated. Because $\Omega_{ij} \not\to \infty$ as $\lambda_i \to \lambda_j$ for distinct eigenvalue derivatives, $\Omega_{ij}(t_2)$ can be written as a Taylor series about t_1:

$$\Omega_{ij}(t_2) = \Omega_{ij}|_{t_1} + \dot{\Omega}_{ij}|_{t_1} \Delta t + \tfrac{1}{2!} \ddot{\Omega}_{ij}|_{t_1} \Delta t^2 + \mathcal{O}(\Delta t^3) \tag{34}$$

The approximation $\underline{\Omega}_{ij}$ is simply a first-order truncation of Eq. (34):

$$\underline{\Omega}_{ij}(t_2) = \Omega_{ij}|_{t_1} + \dot{\Omega}_{ij}|_{t_1} \Delta t \tag{35}$$

Assuming that the third- and higher-order derivatives of Ω_{12} are small compared to $\ddot{\Omega}_{12}$, the largest $\Delta\Omega_{12}$ likely to be encountered would be of the order

$$\Delta\Omega_{12} \simeq \tfrac{1}{2} \ddot{\Omega}_{12}(t_1) \Delta t^2 + \mathcal{O}(\Delta t^3) \tag{36}$$

Using Eqs. (33) and (36), $\Delta\dot{M}$ can be approximated as

$$\Delta\dot{M} \approx \frac{\epsilon^3 \ddot{\Omega}_{12}(t_1)}{(\dot{\lambda}_1 - \dot{\lambda}_2)^2} \left(\boldsymbol{c}_1 \boldsymbol{c}_2^T + \boldsymbol{c}_2 \boldsymbol{c}_1^T \right) \tag{37}$$

which is of order ϵ^3. If there are m pairs of eigenvalues within ϵ of each other, then the total error in $\Delta\dot{M}$ would be the sum of all m corresponding errors, as shown in Eq. (37), and also be of order ϵ^3. □

Reference 1 presents a similar approximation of Ω_{ij} as is shown in Eq. (29). However, there Ω_{ij} is simply set to zero whenever λ_i and λ_j are sensed to be very close numerically. Following a similar error analysis as was just done, it can be shown that setting Ω_{ij} to zero results in an error in $\dot{M}$ of the order of ϵ^2.

Note that the approximation of Ω_{ij} only directly affects $\dot{M}$. The actual error in M in a numerical simulation would be related to the type of numerical integration used. The preceding analysis ignores the error implicit in approximating $\dot{\Omega}_{ij}$. If this approximation is not valid to order ϵ, then an error $\Delta\dot{M}$ would still be of order ϵ^2. This error analysis will always break down whenever the corresponding eigenvalue derivatives lose their distinctness. This situation will be discussed in the following section.

B. Equal Eigenvalue Derivatives

The error $\Delta\dot{M}$ introduced through the approximation in Eq. (29) is more difficult to estimate whenever both $\lambda_i = \lambda_j$ and $\dot{\lambda}_i = \dot{\lambda}_j$. However, it is possible to give a conservative error estimate for $\dot{M}$.

Proposition 3: The continuous, symmetric matrix M is assumed to have a continuous derivative. Let $|\lambda_i - \lambda_j| \leq \epsilon$ and $|\dot{\lambda}_i - \dot{\lambda}_j| \leq \epsilon$ during the time period $[t_1, t_2]$. Then at t_2, where the eigenvalues

start to differ again by more than ϵ, the error in M introduced by the approximation in Eq. (29) will be no more than of order ϵ.

Proof: Without loss of generality, assume that only the first two eigenvectors are close, i.e., $|\lambda_1 - \lambda_2| \leq \epsilon$. The time period $[t_1, t_2]$ spans the time where both eigenvalues and eigenvalue derivatives are almost equal. At worst the λ_1 and λ_2 would differ by ϵ:

$$\lambda_2 = \lambda_1 + \epsilon \tag{38}$$

Equation (23) can be written as

$$M = \lambda_1 E_{12} + \epsilon E_2 + \sum_{i=3}^{n} \lambda_i E_i \tag{39}$$

For this case the approximation in Eq. (29) only affects Ω_{12}. Note that this Ω_{12} term controls the rotational speed of the $\boldsymbol{c}_1$ and $\boldsymbol{c}_2$ eigenvectors in the plane spanned by the two. Equation (39) shows that whenever there are nearly repeated eigenvalues, then the subspace spanned by the corresponding eigenvectors is important, not the individual eigenvectors themselves. Therefore approximating Ω_{12} as in Eq. (29) will lead to an error in M of the order of ϵ as long as $|\lambda_1 - \lambda_2| < \epsilon$. □

Note that, if the eigenvectors remain continuous during $[t_1, t_2]$, then no further modifications need be done after time t_2. However, whenever the eigenvalues and eigenvalue derivatives are equal, it is possible for the eigenvectors to be discontinuous. Let the two integer sets R^- and R^+ be defined as

$$R^- = \{1, 2, \ldots, r\} \tag{40}$$

$$R^+ = \{r+1, r+2, \ldots, n\} \tag{41}$$

Proposition 4: Let M be a continuous symmetric matrix with a continuous derivative $\dot{M}$. Assume M has r repeated eigenvalues, which have equal derivatives. Then if c_i ($i \in R^-$) is discontinuous, then Ω_{ip} is either discontinuous or zero for every $p \in R^+$.

Proof: Using Eqs. (7) and (24), we can express $\dot{E}_i$ as

$$\dot{E}_i = -\sum_{j=1}^{n} \Omega_{ij} \left(\boldsymbol{c}_j \boldsymbol{c}_i^T + \boldsymbol{c}_i \boldsymbol{c}_j^T \right) \qquad i \neq j \tag{42}$$

Without loss of generality, let the first r eigenvalues be repeated. Then E_{1-r} is defined as

$$E_{1-r} = E_1 + \cdots + E_r \tag{43}$$

Because the first r eigenvalues and their derivatives are equal, Eq. (25) can be written as

$$\dot{M} = \dot{\lambda}_1 E_{1-r} + \lambda_1 \dot{E}_{1-r} + \sum_{j=r+1}^{n} \dot{\lambda}_j E_j + \lambda_j \dot{E}_j \tag{44}$$

Because $\dot{M}$ is continuous, the right-hand side of Eq. (44) also would have to be continuous whenever any eigenvectors were discontinuous. By definition the subspace E_{1-r} is invariant to any instantaneous change of base eigenvectors. However, the terms $\dot{E}_{1-r}$ and $\dot{E}_j$ need further investigation. Using Eq. (42), $\dot{E}_{1-r}$ can be reduced to

$$\dot{E}_{1-r} = -\sum_{i=1}^{r} \sum_{j=r+1}^{n} \Omega_{ij} \left(\boldsymbol{c}_j \boldsymbol{c}_i^T + \boldsymbol{c}_i \boldsymbol{c}_j^T \right) \tag{45}$$

where the fact was used that $\Omega_{ji} = -\Omega_{ij}$. As expected, no Ω_{ij} terms relating the eigenvectors of two repeated eigenvalues appear in $\dot{E}_{1-r}$.

The only possibility for $\dot{E}_{1-r}$ to be invariant to any discontinuity in $\boldsymbol{c}_i$ ($i \in R^-$) is for Ω_{ip} to be either discontinuous or zero for every $p \in R^+$. Studying Eq. (42), the same can be said for $\dot{E}_j$. □

Note that the eigenvector derivatives $\dot{\boldsymbol{c}}_i$ ($i \in R^-$) could be anything as long as $\lambda_i = \lambda_j$ and $\dot{\lambda}_i = \dot{\lambda}_j$ ($i, j \in R^-$) (Ref. 6). However, as soon as either the eigenvalues or their derivatives become distinct, the corresponding eigenvectors and derivatives are uniquely determined again. If the eigenvectors are erroneously continuously forward integrated, then the corresponding C matrix would no longer diagonalize the current $M(\boldsymbol{x})$ matrix. What is needed is a method to rotate the two eigenvectors within the plane spanned by them such that they once again diagonalize the mass matrix.

V. Stabilization Using the Jacobi Method

No matter what method we use, and regardless of whether nearly repeated eigenvalues have been encountered, numerical integration errors degrade the accuracy of the eigenvector matrix C. Without further adjustments, the errors will accumulate and C will no longer properly diagonalize the current mass matrix $M(\boldsymbol{x}, t)$. These adjustments are particularly desirable during periods when the Ω_{ij} terms are only being approximated and to handle cases of discontinuous eigenvectors.

We introduce a stabilization method based on the Jacobi method.[13] The Jacobi method has been used for over a century because of its simplicity and stability. It finds the eigenvalues and eigenvectors of a symmetric matrix M by pre- and postmultiplying it by successive orthogonal rotation matrices P_i, as will be briefly outlined.[14]

Let P_i be the ith rotation matrix and Λ_k be the matrix obtained after pre- and postmultiplying M by the first k P_i matrices:

$$P_k^T \cdots P_1^T P_0^T M P_0 P_1 \cdots P_k = \Lambda_k \tag{46}$$

As $k \to \infty$, the matrix Λ_k becomes diagonal with its entries being the eigenvalues of M. The eigenvector matrix C as defined in Eq. (1) is approximated at the kth step as

$$C = P_k^T \cdots P_1^T P_0^T \tag{47}$$

The orthogonal rotation matrix P_{k+1} is defined such that the i, jth entry of Λ_k is zeroed. The matrix P_{k+1} is a diagonal matrix with the ith and jth diagonal elements being $\cos\theta$. The only two nonzero off-diagonal elements are the i, jth element being $-\sin\theta$ and the j, ith element being $\sin\theta$. The rotation angle θ is defined as

$$\theta = \frac{1}{2} \arctan \frac{2\Lambda_{ij}^k}{\Lambda_{ii}^k - \Lambda_{jj}^k} \tag{48}$$

if the diagonal elements are distinct or as

$$\theta = \pi/4 \tag{49}$$

if the diagonal elements are equal. Once the off-diagonal elements are small in magnitude, the procedure is shown to have quadratic convergence.[14] Thus, for the near diagonal mass matrix, we expect very rapid convergence.

In the eigenfactor mass matrix parameterization method presented in this paper, we already have a close approximation $\hat{C}$ of the eigenvector matrix C:

$$\hat{C} M \hat{C}^T = \hat{\Lambda}_0 \tag{50}$$

Because the numerical integration is not perfect, or because some Ω_{ij} terms are only being approximated, the $\hat{\Lambda}_0$ matrix may not be exactly diagonal. To cancel the first off-diagonal element, the rotation matrix P_1 is constructed and operated on $\hat{\Lambda}_0$:

$$P_1^T \hat{\Lambda}_0 P_1 = \hat{\Lambda}_1 \tag{51}$$

This process is then repeated for each of the remaining off-diagonal elements. Assuming the matrix has k significant off-diagonal elements, the adjusted eigenvector matrix $\hat{C}_{\text{ad}}$ becomes

$$\hat{C}_{\text{ad}} = P_k^T \cdots P_1^T \hat{C} \tag{52}$$

Note that each $P_i \hat{C}$ update does not involve a full matrix multiplication. Actually only two eigenvectors are linearly combined, i.e., rotated, to form the new eigenvectors.

If no eigenvector discontinuity had occurred, then it was shown earlier that the error due to approximating Ω_{ij} would be of the order of ϵ^3 for the case of distinct eigenvalue derivatives and at least of

order ϵ for the case of repeated eigenvalue derivatives. In either case, the off-diagonal terms of $\hat{\Lambda}_0$ would be very small before introducing Jacobi rotations. A Jacobi sweep refers to the process of sequentially zeroing each off-diagonal element. Inasmuch as the Jacobi method has quadratic convergence and the off-diagonal terms are of order ϵ^3 or smaller to begin with, one Jacobi sweep will suffice to cancel any diagonalizing error to order ϵ^6 of the C matrix, typically within machine accuracy for that particular integration step.

This Jacobi sweep should be used whenever two eigenvalues are near equal and Ω_{ij} terms are being approximated. Performing the approximations without the Jacobi sweeps was usually found to be stable for a single eigenvalue crossing. However, for repeated cases of near-equal eigenvalue events, the errors accumulate and drive the integration unstable in some simulations. Performing these Jacobi sweeps ensures that the current C and Λ matrices satisfy the necessary condition that they diagonalize the current mass matrix $M(\boldsymbol{x})$ to a high degree of approximation and stabilize the integration process.

For the case of repeated eigenvalues, there is an infinite choice of eigenvectors that will diagonalize M. It was shown that for non-repeated eigenvalue derivatives there is a unique choice of eigenvectors that will generate the proper $\dot{M}$ matrix. The Jacobi method will only correct C such that it diagonalizes M without any regard to consistency with $\dot{M}$. Therefore, small residual state errors can still be expected to accumulate even after the Jacobi sweep. However, numerical studies showed that the Jacobi sweeps dramatically improve the accuracy and the long-term stability of the orthogonal square root eigenfactor parameterization.

If eigenvalues and their derivatives are equal, it was shown earlier that it is possible for the eigenvectors to be discontinuous. During the integration, discontinuous eigenvectors will cause rapidly growing off-diagonal terms of the $\hat{\Lambda}_0$ matrix if no further action is taken. However, performing a Jacobi sweep after each integration step where near-equal eigenvalues occur automatically rotates the two corresponding eigenvectors such that they again diagonalize the current $M(\boldsymbol{x}, t)$ to within machine precision. Therefore, discontinuous eigenvectors are handled with the same Jacobi sweep process used to ensure long-term stability; this effectively restarts the solution whenever a discontinuity is passed.

During the periods where no eigenvalues are close to being equal (the common case), it is not necessary to perform a Jacobi sweep after each integration step. Numerical studies showed that performing it after each integration step caused a slightly larger long-term error than simply performing it periodically due to an increased number of arithmetic operations. At present we have no rigorous means to choose the frequency to perform Jacobi sweeps, but a wide set of admissible choices appears to exist.

VI. Results

A preliminary study of the orthogonal square root eigenfactor parameterization of a mass matrix was performed on the three-link system shown in Fig. 2. The system was chosen to have a low number of degrees of freedom to simplify the analysis of the method. Also, because the system is conservative, the total system energy E remains constant, and this integral of motion provides a convenient way of measuring the current solution error.

All masses and rods are set to unit weight and length, and the system is initially at rest. Initial total energy $E(0) = 1.11355$. The spring constant K is 0.2. The initial orientation angles are $\theta_1 = -90$ deg, $\theta_2 = -30$ deg, and $\theta_3 = 0$ deg. The forward integration is performed with a variable step size Runge–Kutta method. The classic fourth-order Runge–Kutta method is the primary integrator, which is compared to a third-order Runge–Kutta method. If the difference in states between the two integration methods does not lie between two thresholds, then the integration step size is either scaled up or scaled down. In all simulation runs the high-accuracy threshold is always set to 0.015 times the low-accuracy threshold. To compare two eigenvector matrices, the Frobenius norm is used. To compare solutions based on two different integrations, the average integration step size $\bar{h}$ is used. Assuming that the simulation involved n integration steps h_i, then $\bar{h}$ is defined as

$$\bar{h} = \frac{1}{t_f} \sum_{i=0}^{n} h_i^2 \tag{53}$$

The norm of the total energy error is defined as

$$\|\Delta E\|_{L^1(0,t_f;\Re)} = \frac{1}{t_f} \int_0^{t_f} |\Delta E| \, \mathrm{d}t \tag{54}$$

The resulting eigenvalue time evolution for up to 20 s is shown in Fig. 3. The eigenvalues λ_2 and λ_3 become very close at certain times yet are never exactly equal. It appears as though they repel from each other in the last moment and change directions. This phenomenon of eigenvalue near misses is discussed by Arnold in Appendix 10 of Ref. 15. We note that this is but one eigenvalue encounter pathology; for more general problems, the eigenvalues can indeed cross. For the single-parameter system mass matrix $M(t)$, it is very unusual to encounter repeated eigenvalues, but it is not impossible. A two-link manipulator can be configured such that the eigenvalues do periodically become equal. If the manipulator chain is allowed to fork, then it is also possible for eigenvalues to become equal. However, in all physical systems studied so far, no system was found that exhibited crossing eigenvalues λ_i and λ_j and that had nonzero corresponding Ω_{ij} terms. If the eigenvalues actually crossed, then their relative eigenaxis angular velocity was always found to be zero. Although this pattern has been consistently observed, no claims are made that this will occur for all physical systems.

To track the rapid eigenvalue changes near $\lambda_2 \approx \lambda_3$, a variable step size integration method is essential. Every time these two eigenvalues become very close, something interesting happens to corresponding eigenvectors. Away from a near eigenvalue encounter, the eigenvectors oscillate normally, as is seen in their angular velocity measures Ω_{ij} in Fig. 4. However, every close approach of λ_2 and λ_3 causes the corresponding eigenvector axes to switch places. This switching is seen as a clear spike in the Ω_{23} time history in Fig. 4.

The following integration error studies were performed with ϵ set small enough such that Ω_{ij} never invoked the linear approximation of Eq. (29). Just the standard $\dot{S}$ and $\dot{C}$ equations are used. Case 1 is a simulation where no Jacobi sweeps are performed. This is essentially the same as using the method as proposed in Ref. 1. Case 2 is a simulation where a Jacobi sweep is performed after each integration step, and in case 3 the Jacobi sweep is performed only periodically

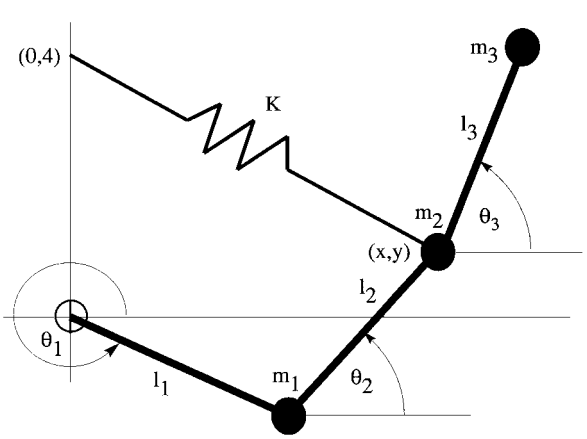

Fig. 2 Three-link manipulator system.

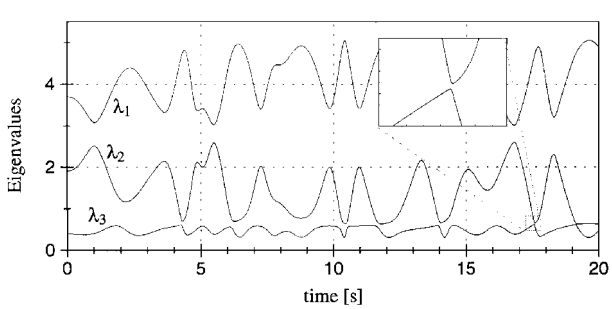

Fig. 3 Mass matrix eigenvalue time history.

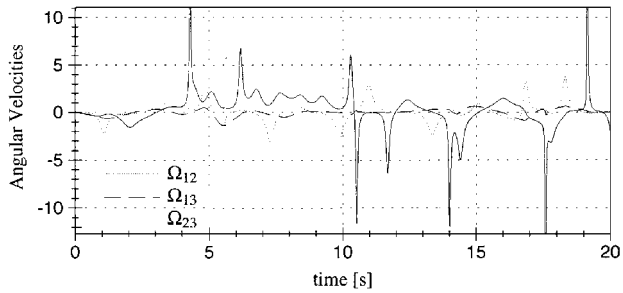

Fig. 4 Eigenvector axis angular velocities.

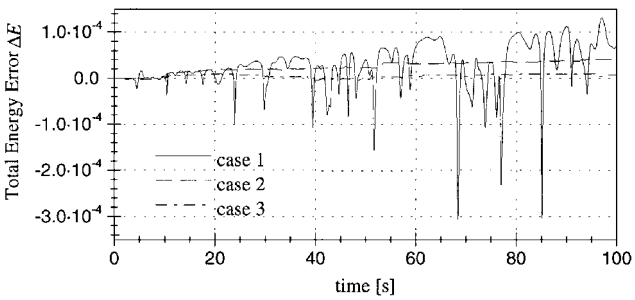

Fig. 5 Total energy error evolutions.

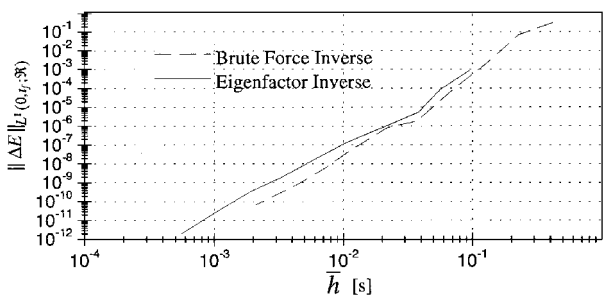

Fig. 6 Computational accuracy comparison.

after each seventh integration step. The time history of the total energy error for the first 100 s is shown in Fig. 5. For each simulation the integration accuracy threshold was 10^{-5}, which resulted in an average step size of about $\bar{h} = 0.04$ s. Without any Jacobi sweeps, the method clearly has a serious long-term stability problem. The energy errors grows more erratic with time. Performing the Jacobi sweep after each integration step clearly smooths out the error curve. However, the additional arithmetic errors are starting to contribute an energy error as well. Performing the Jacobi sweep only periodically after each seventh integration step was found to stabilize and apparently minimize the total energy error. For this case, the errors end up being on the average about one order of magnitude smaller than for the other two cases. This general behavior was found to be true for all integration step sizes and other system configurations tested.

In Fig. 6 the accuracy of the orthogonal square root eigenfactor mass matrix parameterization is compared to the accuracy obtained by performing the inverse M^{-1} directly. For the simple 3×3 system, an exact matrix inverse formulation was used. For a given integration accuracy threshold, the brute force inverse method had one to two orders of magnitude less accuracy. However, this was because the accuracy threshold enforced a larger average integration step size than for the new method. For a fairer comparison, Fig. 6 compares the error norms relative to the average integration step size. Integrating the eigenfactors provides only a small decrease in accuracy compared to the brute force inverse method. The brute force inverse method would be very difficult to beat for this low-order dynamical system because it employs an exact matrix inverse. Taking the inverse for larger systems is known to be computationally very difficult and will inject significant arithmetic error on every time step. However, the integration of the eigenfactors is not expected to be more difficult for higher-order systems. The major integration difficulties are already present in this low-order system, inasmuch as we

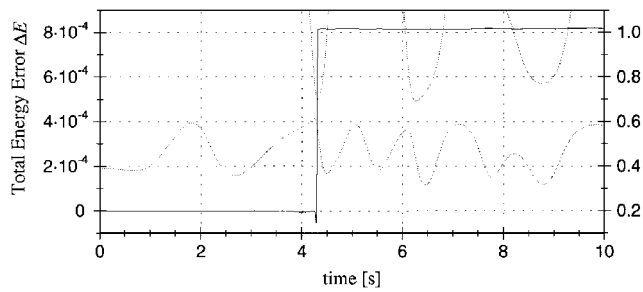

Fig. 7 Total energy error evolution.

have specifically addressed an example with frequent occurrence of near-repeated eigenvalues. Therefore, it is very encouraging to see the eigenfactor method integration error comparable to the exact brute force solution. For higher-dimensioned systems, the eigenfactor method is expected to provide better accuracy and stability. Parallelization of the present algorithm is easily accomplished, and this is conjectured to be an extremely important computational feature for high-dimensional applications. However, this has not been numerically verified and will be studied in future work.

To numerically study the integration accuracy of the Ω_{ij} approximations in Eq. (29), the ϵ term was set to 0.15. The integration accuracy threshold was set to 10^{-5}. Between 4 and 5 s, λ_2 and λ_3 are equal within this bound. The total energy norm for up to 10 s is shown in Fig. 7. The eigenvalues λ_2 and λ_3 are plotted as a comparison in the background with their scale labeled on the right-hand axis. As expected, even with the Jacobi sweeps being performed after any integration steps with Ω_{ij} approximations, there is still a distinct jump in the total energy error when $|\lambda_2 - \lambda_3| < \epsilon$. For the case where the eigenvalues actually become equal and have distinct derivatives, a rough error estimate of ϵ^3 was predicted. Even though the eigenvalues never actually become equal in this situation, the error is still of the order of ϵ^3. The system was also found to be stable for repeated cases of close eigenvalues, even with this crude ϵ threshold. When no Jacobi sweeps were done and the Ω_{ij} were only approximated as a constant, as proposed in Ref. 1, then the solution was found to be unstable for the case of repeatedly close eigenvalues.

Note that the ϵ used in Fig. 7 is much larger than would be normally used with this eigenfactor method. With double precision accuracy, the difference between two eigenvalues can become much smaller before any numerical difficulties occur. Because the linear approximation of Ω_{ij} usually results in errors of the order of ϵ^3, choosing ϵ to be about $1e^{-5}$ would suffice for most applications.

VII. Conclusion

An improved orthogonal square root eigenfactor parameterization of a symmetric, positive definite mass matrix has been presented. A key feature of the proposed integration process is the use of the Jacobi method to stabilize the solution for the eigenfactors. Repeated eigenvalues, repeated eigenvalue derivatives, and discontinuous eigenvectors can be handled by this improved method as illustrated by an example. This method lends itself very well to a massively parallel computer implementation.

References

[1] Oshman, Y., and Bar-Itzhack, I. Y., "Eigenfactor Solution of the Matrix Riccati Equation—A Continuous Square Root Algorithm," *IEEE Transactions on Automatic Control*, Vol. AC-30, No. 10, 1985, pp. 971–978.

[2] Kim, Y., Lee, S., and Junkins, J. L., "Eigenvector Derivatives for Mechanical Second-Order Systems," *Journal of Guidance, Control, and Dynamics*, Vol. 18, No. 4, 1995, pp. 899–906.

[3] Lin, R. M., and Lim, M. K., "Eigenvector Derivatives of Structures with Rigid Body Modes," *AIAA Journal*, Vol. 34, No. 5, 1996, pp. 1083–1085.

[4] Junkins, J. L., and Kim, Y., *Introduction to Dynamics and Control of Flexible Structures*, AIAA Education Series, AIAA, Washington, DC, 1993, pp. 53–57.

[5] Schaub, H., Tsiotras, P., and Junkins, J. L., "Principal Rotation Representations of Proper $N \times N$ Orthogonal Matrices," *International Journal of Engineering Science*, Vol. 33, No. 15, 1995, pp. 2277–2295.

[6] Dailey, R. L., "Eigenvector Derivatives with Repeated Eigenvalues," *AIAA Journal*, Vol. 27, No. 4, 1989, pp. 486–491.

[7]Mills-Curran, W. C., "Calculation of Eigenvector Derivatives for Structures with Repeated Eigenvalues," *AIAA Journal*, Vol. 26, No. 7, 1988, pp. 867–871.

[8]Zhang, D.-W., and Wei, F.-S., "Computation of Eigenvector Derivatives with Repeated Eigenvalues Using a Complete Modal Space," *AIAA Journal*, Vol. 33, No. 9, 1995, pp. 1749–1753.

[9]Zhang, Y.-Q., and Wang, W.-L., "Eigenvector Derivatives of Generalized Nondefective Eigenproblems with Repeated Eigenvalues," *Journal of Engineering for Gas Turbines and Power*, Vol. 117, Jan. 1995, pp. 207–212.

[10]Junkins, J. L., and Turner, J. D., *Optimal Spacecraft Rotational Maneuvers*, Elsevier Science, Amsterdam, 1986, pp. 9–13.

[11]Bar-Itzhack, I. Y., and Markley, F. L., "Minimal Parameter Solution of the Orthogonal Matrix Differential Equation," *IEEE Transactions on Automatic Control*, Vol. 35, No. 3, 1990, pp. 314–317.

[12]Bowen, R. M., and Wang, C.-C., *Introduction to Vectors and Tensors*, Vol. 1, Plenum, New York, 1976, pp. 123–166.

[13]Jacobi, C. G. J., "Über ein leichtes Verfahren die in der Theorie der Säculärstörungen vorkommenden Gleichungen numerisch aufzulösen," *Crelle's Journal*, Vol. 30, 1846, pp. 51–94.

[14]Rutishauser, H., "The Jacobi Method for Real Symmetric Matrices," *Journal of Numerical Methods*, Vol. 9, No. 1, 1966, pp. 1–10.

[15]Arnold, V. I., *Mathematical Methods of Classical Mechanics*, 2nd ed., Springer–Verlag, New York, 1989, pp. 425–437.