

SUPPLEMENTARY MATERIALS: DIAGONAL SCALINGS FOR THE EIGENSTRUCTURE OF ARBITRARY PENCILS*

FROILÁN M. DOPICO[†], MARÍA C. QUINTANA[‡], AND PAUL VAN DOOREN[§]

Abstract. This document contains numerical experiments performed to illustrate the effect of the scaling techniques introduced in the accompanying paper on the accuracy of computed eigenvalues of pencils.

SM1. Introduction. These supplementary materials complete the numerical experiments presented in Section 6 of the accompanying paper. Thus, they should be read after reading that paper since we use the same notations, which are not defined again here for brevity. We often refer to equations and tables in the accompanying paper. These references are easily identified because they do not start with “SM”, in contrast to references to tables or equations in these supplementary materials which all start with “SM”. These materials are organized as follows: Section SM2 presents examples on the accuracy of computed eigenvalues of singular pencils, both square and rectangular, and Section SM3 considers the accuracy of the computed eigenvalues of a pencil arising in the solution of a real-world quadratic eigenvalue problem.

SM2. Numerical examples for singular pencils.

SM2.1. Examples on the accuracy of computed eigenvalues of singular square pencils. In this section, we discuss tests for two families of singular square pencils. The first family includes dense pencils for which the regularization in Section 5 is not needed, while the second one corresponds to sparse pencils for which the regularization is necessary. For completeness, Ward’s method [SM11, SM12] is also considered in the comparisons, because, although it was developed for regular pencils, it has worked on the singular ones of this subsection. As in Subsection 6.2, we generated random singular pencils whose “exact” eigenvalues are known and we used the vectors of chordal distances, $c := \|[c_1, \dots, c_n]\|_2$ for the original pencil $(\lambda B - A)$ (c_{orig}), for the balanced pencil $D_\ell(\lambda B - A)D_r$ constructed by the methods in either Section 3 or 5 (c_{bal}), and for the balanced pencil constructed by Ward’s method (c_{ward}), in order to check the improvements that the different scalings produced on the accuracy of the computed eigenvalues.

The first family of dense pencils considered in this subsection is constructed in

*Supplementary materials for SIMAX MS#M136401.

Funding: This work is part of the “Proyecto de I+D+i PID2019-106362GB-I00 financiado por MCIN/AEI/10.13039/501100011033”. It has been also funded by “Ministerio de Economía, Industria y Competitividad (MINECO)” of Spain and “Fondo Europeo de Desarrollo Regional (FEDER)” of EU through grants MTM2015-65798-P and MTM2017-90682-REDT and by the Madrid Government (Comunidad de Madrid-Spain) under the Multiannual Agreement with UC3M in the line of Excellence of University Professors (EPUC3M23), in the V PRICIT (Regional Programme of Research and Technological Innovation). The research of M. C. Quintana was funded by the “contrato predoctoral” BES-2016-076744 of MINECO and by an Academy of Finland grant (Suomen Akatemian päätös 331240). This work was partially developed while Paul Van Dooren held a “Chair of Excellence UC3M - Banco de Santander” at Universidad Carlos III de Madrid in the academic year 2019-2020.

[†]Departamento de Matemáticas, Universidad Carlos III de Madrid, Avda. Universidad 30, 28911 Leganés, Spain (dopico@math.uc3m.es).

[‡]Department of Mathematics and Systems Analysis, Aalto University, Otakaari 1, Espoo, Finland (maria.quintanaponce@aalto.fi).

[§]Department of Mathematical Engineering, Université catholique de Louvain, Avenue Georges Lemaître 4, B-1348 Louvain-la-Neuve, Belgium (paul.vandooren@uclouvain.be).

the same way as the pencils in Table 9, but we replaced one of the diagonal pairs of the 500×500 pencil $(\lambda\Lambda_B - \Lambda_A)$ generated in the regular example by two zeros, thus creating a singular pencil. Each transformed pencil $(\lambda B - A) := T_\ell(\lambda\Lambda_B - \Lambda_A)T_r$ is therefore also singular, but its left and right rational null spaces are both of dimension 1 and their minimal bases are formed by constant vectors [SM10]. For that reason, the regular part of that singular pencil has dimension 499×499 and its eigenvalues are the remaining 499 eigenvalues of $(\lambda\Lambda_B - \Lambda_A)$. If we follow the same procedure as in the regular experiment, the QZ -algorithm applied to $(\lambda B - A)$ should in principle yield arbitrary eigenvalues, since it is known that the QZ -algorithm is backward stable and that there exist arbitrarily small perturbations of square singular pencils that make them regular, but with arbitrary spectrum in the complex plane [SM10]. However, it has been shown that such perturbations are very particular, and that, generically, tiny perturbations of a singular square pencil makes it regular with eigenvalues that are tiny perturbations of the eigenvalues of the unperturbed singular pencil, together with some other “arbitrary” eigenvalues determined by the perturbation [SM2, SM3]. Even more, starting from these ideas, it has been shown very recently that it is possible to define sensible and useful “weak” condition numbers for the eigenvalues of a singular square pencil [SM8]. This explains the well-known fact that, in practice, the QZ -algorithm applied to a singular square matrix pencil finds almost always its eigenvalues, albeit with some loss of accuracy. Therefore, it makes sense to apply the QZ algorithm to our generated singular pencils as well as to their scaled versions. The numerical results are reported in Table SM1, where each row corresponds to a value of k taken in increasing order from $k = 1 : 5 : 41$ as in Table 9. We generated the data just as in the experiment for regular pencils in Table 9, except for the one eigenvalue replaced by 0/0 or, in other words, by NaN. When comparing the “original” spectrum with the computed one, we excluded NaN in the original set and looked for the best matching 499 eigenvalues in the “computed” spectrum. It is clear from Table SM1 that the balancing proposed in Section 3 also improves the accuracy of the computed eigenvalues of singular square pencils, both with respect to the original pencil and with respect to the one balanced by Ward’s method, and that needs a small number of steps to converge.

Though the direct use of the QZ -algorithm is a simple option for computing the eigenvalues of a singular square pencil when the accuracy requirements are moderate, the correct handling of a singular pencil is to first “deflate” its left and right null spaces, and then compute the spectrum of the regular part of that singular pencil, i.e., to apply the staircase algorithm (see [SM10]). In this experiment, it turns out that the left and right null spaces are one-dimensional and are given, respectively, by the left null vector of $[A \ B]$, and by the right null vector of $\begin{bmatrix} A \\ B \end{bmatrix}$, which we both computed using a singular value decomposition of these compound matrices. After this deflation was applied to the original pencil $(\lambda B - A)$, to the pencil $D_\ell(\lambda B - A)D_r$ scaled by the method in Section 3 and to the one balanced by Ward’s method, we again computed the spectrum of the deflated pencils with the QZ -algorithm. The results for the same data as reported in Table SM1 are now reported in Table SM2. The results in this case are similar in both tables. We also added three columns with the sensitivities of the deflation in the original pencil γ_{orig} and in the balanced ones by the method in Section 3 and Ward’s method, γ_{bal} and γ_{ward} . We measured the sensitivity of the left and right null vectors defining the deflation of a singular pencil $\lambda B - A$, by

$$(SM2.1) \quad \gamma := \max\left(\frac{\sigma_n \begin{bmatrix} A \\ B \end{bmatrix}}{\sigma_{n-1} \begin{bmatrix} A \\ B \end{bmatrix}}, \frac{\sigma_n [A \ B]}{\sigma_{n-1} [A \ B]}\right),$$

TABLE SM1

Eigenvalue accuracy of the plain QZ-algorithm for singular 500×500 dense pencils: for the original pencil, for the pencil balanced by applying the algorithm in Appendix A with $r = c = n\mathbf{1}_n$ and $\text{tol}=1$ to $M = |A|^{\circ 2} + |B|^{\circ 2}$, and for the pencil balanced by Ward's method. The improvement in the scaling of M produced by the algorithm in Appendix A is also shown in terms of $q_S(M_{orig})$ and $q_S(M_{scal})$ (see (5.6)), as well as the number of its steps until convergence

$\kappa(T_\ell)$	$\kappa(T_r)$	C_{orig}	C_{bal}	C_{ward}	C_{bal}/C_{orig}	C_{bal}/C_{ward}
4.30e+03	4.10e+03	1.88e-12	1.88e-12	8.27e-12	1.00e+00	2.28e-01
1.69e+04	2.12e+04	1.77e-11	1.85e-12	6.17e-12	1.04e-01	2.99e-01
1.06e+06	9.83e+04	1.88e-11	1.19e-11	5.04e-12	6.34e-01	2.37e+00
7.47e+05	2.73e+06	1.98e-10	1.40e-10	7.13e-11	7.08e-01	1.97e+00
1.20e+08	6.49e+08	1.62e-08	4.13e-11	4.13e-09	2.55e-03	9.99e-03
2.32e+10	2.75e+09	5.20e-07	5.00e-09	2.15e-07	9.62e-03	2.33e-02
3.59e+13	2.59e+12	3.25e-03	2.83e-07	5.40e-05	8.71e-05	5.24e-03
1.63e+16	3.03e+13	3.46e-02	3.55e-05	3.84e-03	1.03e-03	9.25e-03
1.63e+18	1.48e+14	8.15e-02	9.12e-06	1.22e-02	1.12e-04	7.46e-04

$q_S(M_{orig})$	$q_S(M_{scal})$	steps
1.57e+00	1.57e+00	1
1.09e+03	6.04e+00	3
1.40e+06	8.91e+00	7
2.66e+09	9.43e+00	8
1.10e+12	1.01e+01	13
1.31e+16	9.26e+00	13
1.13e+20	1.43e+01	16
1.72e+25	1.20e+01	17
2.11e+26	1.16e+01	18

i.e. the largest ratio between the two smallest singular values of the matrices that define these null vectors. It is an indication about how much these vectors can rotate when perturbing the pencil. It is easy to see from the data that the accuracy of the computed eigenvalues of the deflated pencil is closely related to the sensitivity of the deflation itself.

TABLE SM2

Eigenvalue accuracy of the staircase algorithm for exactly the same singular 500×500 dense pencils of Table SM1

C_{orig}	C_{bal}	C_{ward}	C_{bal}/C_{orig}	C_{bal}/C_{ward}	γ_{orig}	γ_{bal}	γ_{ward}
2.23e-13	2.23e-13	2.33e-13	1.0e+00	9.57e-01	5.79e-13	5.79e-13	6.59e-13
4.53e-13	4.68e-13	2.89e-13	1.03e+00	1.62e+00	1.48e-11	4.96e-12	7.49e-12
6.92e-13	9.11e-13	2.00e-12	1.32e+00	4.56e-01	2.37e-10	2.33e-12	1.17e-10
8.36e-11	1.63e-11	9.76e-12	1.95e-01	1.67e+00	1.33e-07	5.10e-11	4.84e-08
6.49e-10	1.12e-11	8.46e-11	1.73e-02	1.33e-01	1.24e-06	2.02e-11	1.17e-07
1.41e-07	5.06e-09	2.03e-07	3.59e-02	2.49e-02	1.53e-03	2.02e-09	1.35e-04
7.22e-04	1.42e-06	1.03e-06	1.96e-03	1.38e+00	9.62e-01	3.43e-07	1.99e-01
3.33e-02	1.25e-06	9.17e-03	3.76e-05	1.36e-04	2.51e-01	2.18e-06	5.24e-01
1.08e-01	4.31e-07	1.84e-03	3.97e-06	2.34e-04	3.87e-01	4.59e-07	7.27e-01

The second family of sparse singular pencils considered in this subsection is a family of 400×400 permuted block diagonal pencils generated as follows. Set, for

simplicity, $m_1 = 140$ and $n_1 = 260$. Then

$$(SM2.2) \quad \lambda B - A := P \begin{bmatrix} \lambda B_1 - A_1 & \\ & \lambda B_2 - A_2 \end{bmatrix} Q,$$

with P, Q random 400×400 permutation matrices and

$$\begin{aligned} \lambda B_1 - A_1 &= T_{\ell 1} \begin{bmatrix} \lambda \Lambda_{B1} - \Lambda_{A1} & \\ & 0_{1 \times (n_1 - m_1 + 1)} \end{bmatrix} T_{r1}, \\ \lambda B_2 - A_2 &= T_{\ell 2} \begin{bmatrix} \lambda \Lambda_{B2} - \Lambda_{A2} & \\ & 0_{(n_1 - m_1 + 1) \times 1} \end{bmatrix} T_{r2}, \end{aligned}$$

where $\lambda \Lambda_{B1} - \Lambda_{A1}, \lambda \Lambda_{B2} - \Lambda_{A2}$ are random $(m_1 - 1) \times (m_1 - 1)$ diagonal regular pencils in standard normal form [SM7] which contain the “exact” eigenvalues of $\lambda B - A$, and the entries of $T_{\ell 1} \in \mathbb{R}^{m_1 \times m_1}, T_{r2} \in \mathbb{R}^{m_1 \times m_1}, T_{\ell 2} \in \mathbb{R}^{n_1 \times n_1}, T_{r1} \in \mathbb{R}^{n_1 \times n_1}$ are k th powers of normally distributed random numbers, for $k = 1 : 5 : 41$. Observe that the normal rank [SM10] of these pencils is $rg = 2(m_1 - 1) = 278$, that their left and right rational null spaces are both of dimension 122 and that their minimal bases are formed by constant vectors. This mean that they are given again, respectively, by the left null vectors of $[A \ B]$, and by the right null vectors of $\begin{bmatrix} A \\ B \end{bmatrix}$, which were computed again using a singular value decomposition of these compound matrices. This allowed us to deflate these right and left null spaces and to obtain the regular parts of such pencils by multiplying $\lambda B - A$ on the left by the rg left singular vectors of $[A \ B]$ corresponding to its rg largest singular values and on the right by the rg right singular vectors of $\begin{bmatrix} A \\ B \end{bmatrix}$ corresponding to its rg largest singular values. The application of the QZ algorithm to these regular parts yielded the eigenvalues of these highly singular pencils and we did it for the original pencil $(\lambda B - A)$, for the pencil $D_\ell(\lambda B - A)D_r$ scaled by the *regularized* method in Section 5 and for the one balanced by Ward’s method. The plain QZ algorithm can also be applied directly to the pencils in (SM2.2), but it produces much larger errors than the staircase algorithm described above due to the high singularity of these pencils. The results for the staircase algorithm are shown in Table SM3, where each row corresponds to a value of k , and are discussed in the next paragraph.

The matrices M corresponding to the pencils in (SM2.2) are very far from having total support and the Sinkhorn-Knopp algorithm applied to them with $\mathbf{tol} = 1$ did not converge because it produced diagonal matrices D_ℓ, D_r with zero diagonal entries due to underflows. Then, we regularized the problem by applying the algorithm in Appendix A with $r = c = (2n)\mathbf{1}_{2n}$ and $\mathbf{tol} = 1$ to M_α^{c2} in (5.3) with $\alpha = 0.5$. Observe, that this yielded factors $q_S(M_{scal})$ very far from 1 but much smaller than the factors of the original matrices $q_S(M_{orig})$. Interestingly, the factors $q_S(M_{scal})$ did not improve by taking much smaller values of α . Despite this fact, the impact of the regularized scaling on the accuracy of the computed eigenvalues is impressive both in comparison with the original pencils and with the pencils scaled by Ward’s method. The new regularized method leads to the computation of very accurate eigenvalues in a problem which is extremely difficult in terms of the high singularity and of the high unbalancing of the considered pencils. We do not know any other method in the literature that can achieve such results. Moreover, the numbers of steps until convergence are still moderate taking into account the sparsity and the strong unbalancing of the pencils, and make the cost of the scaling considerably smaller than the cost of computing the eigenvalues. Finally note that Table SM3 also includes the sensitivities of the deflations $\gamma_{orig}, \gamma_{bal}$ and γ_{ward} as in Table SM2. They were computed replacing $n - 1$

TABLE SM3

Eigenvalue accuracy of the staircase algorithm for singular 400×400 sparse pencils: for the original pencil, for the pencil balanced by applying the algorithm in Appendix A with $r = c = (2n)\mathbf{1}_{2n}$ and $\text{tol} = 1$ to $M_\alpha^{\circ 2}$ in (5.3) with $\alpha = 0.5$, and for the pencil balanced by Ward's method. The improvement in the scaling of $M = |A|^{\circ 2} + |B|^{\circ 2}$ produced by the algorithm in Appendix A applied to $M_\alpha^{\circ 2}$ is also shown in terms of $q_S(M_{orig})$ and $q_S(M_{scal})$, as well as the number of its steps until convergence. The last column of the second table shows that the plain QZ-algorithm produces much larger errors for these pencils. For brevity this is only shown for the pencils balanced by the algorithm in Appendix A, but the same happens for the other pencils

C_{orig}	C_{bal}	C_{ward}	C_{bal}/C_{orig}	C_{bal}/C_{ward}	γ_{orig}	γ_{bal}	γ_{ward}
1.98e-14	2.25e-14	2.15e-14	1.14e+00	1.05e+00	1.10e-13	1.27e-13	9.51e-14
3.13e-14	2.10e-14	2.39e-14	6.71e-01	8.80e-01	4.29e-12	3.29e-13	1.38e-12
3.40e-12	4.49e-14	2.72e-13	1.32e-02	1.65e-01	3.80e-10	1.28e-12	5.64e-11
1.76e-11	4.76e-13	2.69e-12	2.71e-02	1.77e-01	3.70e-07	5.02e-11	1.01e-07
3.17e-08	9.47e-13	4.79e-11	2.99e-05	1.98e-02	2.26e-04	2.52e-10	1.90e-07
7.84e-03	7.43e-11	1.10e-08	9.48e-09	6.74e-03	1.0e+00	1.20e-09	1.11e-03
2.31e-04	1.21e-10	5.74e-07	5.23e-07	2.11e-04	1.0e+00	5.42e-08	1.34e-02
1.93e-02	3.32e-08	2.73e-02	1.72e-06	1.22e-06	1.0e+00	2.55e-06	1.0e+00
6.46e-01	4.64e-10	4.19e-03	7.17e-10	1.11e-07	1.0e+00	1.21e-07	1.0e+00

$q_S(M_{orig})$	$q_S(M_{scal})$	steps	C_{bal} plain QZ
3.99e+00	9.08e+00	16	8.56e-07
9.51e+04	8.97e+01	30	8.93e-07
1.75e+09	2.85e+03	45	8.40e-07
5.84e+13	1.82e+05	66	7.01e-07
1.69e+17	5.18e+05	81	2.44e-07
5.26e+23	1.41e+07	100	2.57e-07
1.12e+22	7.74e+06	112	9.03e-08
1.49e+26	1.74e+09	130	5.02e-02
2.75e+36	1.58e+12	149	5.96e-03

and n in (SM2.1) by rg and $rg + 1$, respectively, where $rg = 278$ is the normal rank of the pencils. We also observe in Table SM3 a strong relation between the errors in the eigenvalues and the deflation sensitivities.

The experiments in this section show that the balancing procedures of the paper accompanying these supplementary materials improve the accuracy of the eigenvalue computation of square singular pencils as well as the sensitivity of the deflation of the regular part of a singular pencil. We briefly mention that recently an alternative robust method to the staircase algorithm has been proposed for computing the eigenvalues of singular pencils [SM4]. This new method is related to the ideas in [SM2, SM3, SM8] and its accuracy should also improve by using our scaling strategies.

SM2.2. Examples on the accuracy of computed eigenvalues of rectangular pencils. In this section we discuss briefly tests for two families of rectangular pencils that are related to the families in Subsection SM2.1. The first family includes dense pencils for which the regularization in Section 5 is not needed, while the second one corresponds to sparse pencils for which the regularization is necessary. Ward's method is not considered since it does not work for rectangular pencils. All the considered pencils $\lambda B - A$ have the minimal bases of their left and right null spaces formed by constant vectors. Thus, the computation of their eigenvalues is performed via the variant of the staircase algorithm described in the previous subsection, i.e., computing first the regular parts of these pencils with the singular value decompositions of the

compound matrices $[A \ B]$ and $\begin{bmatrix} A \\ B \end{bmatrix}$, and then applying the QZ -algorithm to the regular parts. We use the same notation and test magnitudes as in Subsection SM2.1.

In the first family of tests of this subsection, we generated 150×450 random pencils of the form $\lambda B - A = T_\ell \text{diag}(\lambda \Lambda_B - \Lambda_A, 0_{1 \times 301}) T_r$, where $(\lambda \Lambda_B - \Lambda_A)$ is diagonal, is in standard normal form [SM7], has dimension 149×149 and contains the “exact” eigenvalues of $\lambda B - A$. The elements of the random square nonsingular matrices $T_\ell \in \mathbb{R}^{150 \times 150}$ and $T_r \in \mathbb{R}^{450 \times 450}$ are k th powers of normally distributed random numbers for $k = 1 : 5 : 41$. These pencils are dense and then the regularization in Subsection 5.1 was not needed. The results are shown in Table SM4 (each row corresponds to a value of k) and illustrate the very positive effect of the scaling technique of Section 4 on the accuracy of computed eigenvalues and its low computational cost.

TABLE SM4

Eigenvalue accuracy of the staircase algorithm for rectangular 150×450 dense pencils: for the original pencil and for the pencil balanced by applying the algorithm in Appendix A with $r = n\mathbf{1}_m$, $c = m\mathbf{1}_n$ and $\text{tol} = 1$ to $M = |A|^{\circ 2} + |B|^{\circ 2}$. The improvement in the scaling produced by the algorithm in Appendix A is also shown in terms of $q_S(M_{\text{orig}})$ and $q_S(M_{\text{scal}})$, as well as the number of its steps until convergence.

C_{orig}	C_{bal}	$C_{\text{bal}}/C_{\text{orig}}$	γ_{orig}	γ_{bal}	$q_S(M_{\text{orig}})$	$q_S(M_{\text{scal}})$	steps
9.96e-15	9.96e-15	1.00e+00	1.01e-13	1.01e-13	2.29e+00	2.29e+00	2
1.95e-14	1.08e-14	5.52e-01	7.97e-13	1.97e-13	4.94e+03	7.77e+00	4
2.62e-13	1.06e-14	4.03e-02	3.03e-10	1.57e-13	1.22e+08	9.66e+00	7
2.27e-12	1.29e-14	5.68e-03	1.31e-08	7.73e-13	4.32e+11	1.06e+01	9
5.61e-09	1.97e-13	3.52e-05	1.39e-04	1.72e-11	1.36e+16	1.17e+01	12
1.51e-05	1.20e-13	7.97e-09	1.95e-01	5.78e-12	8.19e+23	1.07e+01	14
6.03e-05	1.08e-12	1.79e-08	8.08e-03	9.12e-12	3.51e+22	1.27e+01	21
5.49e-02	1.72e-11	3.13e-10	1.00e+00	1.36e-09	2.39e+29	1.17e+01	16
9.76e-02	8.40e-12	8.60e-11	1.00e+00	8.24e-10	1.24e+31	1.32e+01	24

TABLE SM5

Eigenvalue accuracy of the staircase algorithm for singular 700×450 sparse pencils: for the original pencil and for the pencil balanced by applying the algorithm in Appendix A with $r = c = v$ in (5.8) and $\text{tol} = 1$ to $M_\alpha^{\circ 2}$ in (5.3) with $\alpha = 0.5$. The improvement in the scaling of $M = |A|^{\circ 2} + |B|^{\circ 2}$ produced by the algorithm in Appendix A applied to $M_\alpha^{\circ 2}$ is also shown in terms of $q_S(M_{\text{orig}})$ and $q_S(M_{\text{scal}})$, as well as the number of its steps until convergence.

C_{orig}	C_{bal}	$C_{\text{bal}}/C_{\text{orig}}$	γ_{orig}	γ_{bal}	$q_S(M_{\text{orig}})$	$q_S(M_{\text{scal}})$	steps
1.43e-14	1.26e-14	8.86e-01	8.91e-14	8.70e-14	5.54e+01	3.27e+01	7
1.73e-14	1.39e-14	8.06e-01	4.05e-12	1.25e-13	4.64e+06	9.30e+03	13
2.81e-13	3.75e-14	1.34e-01	2.74e-10	1.30e-12	3.10e+11	1.80e+06	26
1.77e-11	1.98e-14	1.12e-03	3.28e-08	4.72e-12	5.14e+19	1.42e+10	32
2.42e-06	6.23e-14	2.58e-08	1.81e-03	1.27e-11	5.87e+28	1.09e+13	46
2.42e-02	1.15e-10	4.77e-09	1.00e+00	1.85e-08	4.53e+29	1.11e+18	46
1.69e-04	2.24e-11	1.32e-07	9.84e-01	1.07e-07	6.95e+37	1.42e+20	68
4.10e-03	2.83e-11	6.88e-09	1.00e+00	4.18e-06	9.32e+39	4.30e+22	84
9.91e-01	6.07e-11	6.13e-11	1.00e+00	1.03e-07	2.72e+44	9.90e+22	87

For describing the second considered family of sparse rectangular pencils, we need the parameters $m_1 = 100$, $n_1 = 400$, $m_2 = 600$ and $n_2 = 50$. Then, the pencils have the structure of those in (SM2.2) but with the following changes in $\lambda B_2 - A_2$: the dimension of $\lambda \Lambda_{B_2} - \Lambda_{A_2}$ becomes $(n_2 - 1) \times (n_2 - 1)$ and $0_{(n_1 - m_1 + 1) \times 1}$ is replaced by $0_{(m_2 - n_2 + 1) \times 1}$. This implies that $T_{\ell 2} \in \mathbb{R}^{m_2 \times m_2}$ and $T_{r 2} \in \mathbb{R}^{n_2 \times n_2}$. For these pencils the algorithm in Appendix A with $r = n\mathbf{1}_m$, $c = m\mathbf{1}_n$ and $\text{tol} = 1$ applied to M did

not converge and we used the scaling described in Subsection 5.1 with $\alpha = 0.5$. The results are shown in Table SM5 (each row corresponds to a value of $k = 1 : 5 : 41$) and illustrate again the impressive positive effect of the new scaling technique on the accuracy of computed eigenvalues and its low computational cost. The values of $q_S(M_{scal})$ did not improve by considering very small values of α .

SM3. A numerical example related to the dynamic behavior of a nuclear power plant. In this last numerical example, we consider the regular quadratic polynomial eigenvalue problem `power_plant` included in the MATLAB toolbox [SM1]. This problem is defined by an 8×8 quadratic matrix polynomial $Q(\lambda) = \lambda^2 M + \lambda D + K$ describing a simple model for the dynamic behavior of a nuclear plant [SM6, SM9]. The mass matrix $M \in \mathbb{R}^{8 \times 8}$ and the damping matrix $D \in \mathbb{R}^{8 \times 8}$ are real symmetric and the stiffness matrix is of the form $K = (1 + i\mu)K_0$, with $K_0 = K_0^T \in \mathbb{R}^{8 \times 8}$, i the imaginary unit and μ a real parameter describing the hysteretic damping of the problem. The matrices M, D and K are badly scaled.

The standard way of solving quadratic eigenvalue problems is via linearizations, i.e., constructing from the coefficients of $Q(\lambda) = \lambda^2 M + \lambda D + K$ a matrix pencil that has exactly the same eigenvalues as the polynomial [SM5, SM9] and, then, applying the QZ or the staircase algorithm to the linearization. The most popular linearization is the first Frobenius companion form, which is defined as follows:

$$(SM3.1) \quad \lambda B - A = \begin{bmatrix} \lambda M + D & K \\ -I & \lambda I \end{bmatrix}.$$

Another popular linearization is the one used by the command of MATLAB for computing the eigenvalues of a matrix polynomial, `polyeig`, which is given by

$$(SM3.2) \quad \lambda B_{MAT} - A_{MAT} = \begin{bmatrix} \lambda D + K & \lambda M \\ -\lambda I & I \end{bmatrix}.$$

In this example, we consider four instances of the `power_plant` quadratic polynomial for the values of the parameter $\mu = 0.2, 0.5, 0.8, 1.1$ (0.2 is the default value in the toolbox). For each of these quadratic matrix polynomials, we take as “exact” eigenvalues those computed with the variable precision arithmetic facility of MATLAB `vpa` with 64 decimal digits, which corresponds to a unit roundoff $\approx 10^{-64}$. Then, we compute in standard double precision the eigenvalues of $Q(\lambda)$ by three methods and compute the norm c of their vectors of chordal distances to the “exact” eigenvalues. These methods are: (a) the QZ-algorithm applied to the pencil $\lambda B - A$ in (SM3.1) (c_{orig}); (b) the QZ-algorithm applied to the scaled pencil $D_\ell(\lambda B - A)D_r$ (c_{bal}) obtained by applying the algorithm in Appendix A with $r = c = n\mathbf{1}_n$ to $M = |A|^{\circ 2} + |B|^{\circ 2}$, i.e., without regularization; (c) the `polyeig` command of MATLAB, which amounts to QZ applied to (SM3.2) (c_{MAT}). The results are shown in Table SM6, where each line corresponds to one value of $\mu = 0.2, 0.5, 0.8, 1.1$. The improvement in accuracy obtained via the scaling in Appendix A applied to M is spectacular, yielding for this highly unbalanced real-world pencil eigenvalues with full machine accuracy, in contrast with the poor accuracy obtained by working directly on the unscaled linearizations. Moreover, the convergence of the algorithm is very fast due to the use of the relaxed stopping criterion `tol = 1`, taking just 5 iterations.

This example illustrates that the commands `polyeig` and `eig(A,B)` of MATLAB would greatly benefit of including by default a scaling technique, as it is done in the command `eig(A)`.

TABLE SM6

Eigenvalue accuracies for four instances of the `power_plant` quadratic matrix polynomial. The eigenvalues were computed via (a) the original Frobenius companion form (SM3.1), (b) the Frobenius companion form balanced by applying the algorithm in Appendix A with $r = c = n\mathbf{1}_n$ and $\text{tol} = 1$ to $M = |A|^{\circ 2} + |B|^{\circ 2}$, and (c) the command `polyeig` of MATLAB. The improvement in the scaling of M produced by the algorithm in Appendix A is also shown in terms of $q_S(M_{orig})$ and $q_S(M_{scal})$, as well as the number of its steps until convergence.

C_{orig}	C_{bal}	C_{MAT}	C_{bal}/C_{orig}	C_{bal}/C_{MAT}	$q_S(M_{orig})$	$q_S(M_{scal})$	steps
1.5e-05	1.8e-16	1.6e-06	1.2e-11	1.1e-10	7.6e+25	8.7	5
1.6e-05	1.1e-16	4.1e-05	6.6e-12	2.6e-12	9.2e+25	8.7	5
2.0e-04	2.0e-16	1.0e-05	1.0e-12	2.0e-11	1.2e+26	4.6	5
6.3e-05	1.3e-16	1.9e-06	2.1e-12	6.9e-11	1.6e+26	9.6	5

REFERENCES

- [SM1] T. Betcke, N. J. Higham, V. Mehrmann, C. Schröder and F. Tisseur, NLEVP: A collection of nonlinear eigenvalue problems, *ACM Trans. Math. Software*, 39(2) (2013) 1-28.
- [SM2] F. De Terán and F. M. Dopico, First order spectral perturbation theory of square singular matrix polynomials, *Linear Algebra Appl.*, 432(4) (2010) 892–910.
- [SM3] F. De Terán, F. M. Dopico and J. Moro, First order spectral perturbation theory of square singular matrix pencils, *Linear Algebra Appl.*, 429 (2008) 548-576.
- [SM4] M. E. Hochstenbach, C. Mehl, and B. Plestenjak, Solving singular generalized eigenvalue problems by a rank-completing perturbation, *SIAM J. Matrix Anal. Appl.*, 40(3) (2019) 1022-1046.
- [SM5] I. Gohberg, P. Lancaster and L. Rodman, *Matrix Polynomials*, Academic Press, New York, 1982.
- [SM6] T. Itoh, Damped vibration mode superposition method for dynamic response analysis, *Earthquake Engrg. Struct. Dyn.*, 2 (1973) 47–57.
- [SM7] D. Lemonnier and P. Van Dooren, Balancing regular matrix pencils, *SIAM J. Matrix Anal. Appl.*, 28(1) (2006) 253–263.
- [SM8] M. Lotz and V. Noferini, Wilkinson’s bus: Weak condition numbers, with an application to singular polynomial eigenproblems, *Found. Comput. Math.*, 20 (2020) 1439-1473.
- [SM9] F. Tisseur and K. Meerbergen, The quadratic eigenvalue problem, *SIAM Rev.* 43(2) (2001) 235–286.
- [SM10] P. Van Dooren, The computation of Kronecker’s canonical form of a singular pencil, *Linear Algebra Appl.*, 27 (1979) 103–141.
- [SM11] R. C. Ward, Balancing the generalized eigenvalue problem, *SIAM J. Sci. Stat. Comput.*, 2 (1981) 141–152.
- [SM12] M. Weder, `BALANCE2` Balancing generalized eigenvalue problem, (<https://www.mathworks.com/matlabcentral/fileexchange/49719-balance2-balancing-generalized-eigenvalue-problem>), MATLAB Central File Exchange, 2015 (retrieved May 5, 2021).